arXiv:1410.1776v1 [cs.AI] 7 Oct 2014

Ontology-based Representation and Reasoning on Process Models: A Logic Programming Approach

Fabrizio Smith and Maurizio Proietti
National Research Council
Istituto di Analisi dei Sistemi ed Informatica "Antonio Ruberti"
Via dei Taurini 19, 00185, Roma, Italy
{fabrizio.smith, maurizio.proietti}@iasi.cnr.it

Abstract

We propose a framework grounded in Logic Programming for representing and reasoning about business processes from both the procedural and ontological point of views. In particular, our goal is threefold: (1) define a logical language and a formal semantics for process models enriched with ontology-based annotations; (2) provide an effective inference mechanism that supports the combination of reasoning services dealing with the structural definition of a process model, its behavior, and the domain knowledge related to the participating business entities; (3) implement such a theoretical framework into a process modeling and reasoning platform. To this end we define a process ontology coping with a relevant fragment of the popular BPMN modeling notation. The behavioral semantics of a process is defined as a state transition system by following an approach similar to the Fluent Calculus, and allows us to specify state change in terms of preconditions and effects of the enactment of activities. Then we show how the procedural process knowledge can be seamlessly integrated with the domain knowledge specified by using the OWL 2 RL rule-based ontology language. Our framework provides a wide range of reasoning services, including CTL model checking, which can be performed by using standard Logic Programming inference engines through a goal-oriented, efficient, sound and complete evaluation procedure. We also present a software environment implementing the proposed framework, and we report on an experimental evaluation of the system, whose results are encouraging and show the viability of the approach.

Keywords: Business Processes, Ontologies, Logic Programming, Knowledge Representation, Verification.

1 Introduction

The adoption of structured and systematic approaches for the management of Business Processes (BPs) that operate within an organization is constantly gaining popularity, especially in medium to large organizations such as manufacturing enterprises, service providers, and public administrations. The core of such approaches is the development of BP models that represent the knowledge about processes in machine accessible form. One of the main advantages of process modeling is that it enables automated analysis facilities, such as the verification that the requirements specified over the models are enforced. The automated analysis issue is addressed in the BP Management (BPM) community mainly from a control flow perspective, with the aim of verifying whether the behavior of the modeled system presents logical errors (see, for instance, the notion of soundness [64]).

Unfortunately, standard BP modeling languages are not fully adequate to capture process knowledge in all its aspects. While their focus is on the procedural representation of a BP as a workflow graph that specifies the planned order of operations, the domain knowledge regarding the entities involved in such a process, i.e., the business environment in which processes are carried out, is often left implicit. This kind of knowledge is typically expressed through natural language comments and labels attached to the models, which constitute very limited, informal and ambiguous pieces of information. The lack of a formal representation of the domain knowledge within process models is widely recognized as an obstacle for the further automation of BPM tools and methodologies that effectively support process analysis, retrieval, and reuse [32].

In order to overcome this limitation, the application of well-established techniques stemming from the area of Knowledge Representation in the domains of BP modeling [32, 17, 36, 66] and Web Services [12, 22] has been shown to be a promising approach. In particular, the use of computational ontologies is the most established approach for representing in a machine processable way the knowledge about the domain where business processes operate, providing formal definitions for the basic entities involved in a process, such as activities, actors, data items, and the relations between them. However, there are still several open issues regarding the combination of BP modeling languages (with their execution semantics) and ontologies, and the accomplishment of behavioral reasoning tasks involving both these components. Indeed, most of the approaches developed for the semantic enrichment of process models or Web Services (such as the above cited ones) do not provide an adequate model theory nor an axiomatization to capture and reasoning on dynamic aspects of process descriptions. On the other hand, approaches based on action languages developed in AI (e.g., [57, 6, 44]) are very expressive formalisms that can be used to simultaneously capture the process and the domain knowledge, but they are too general to be applied to BP modeling, and must be suitably restricted not only towards decidability of reasoning but also to reflect the peculiarities of processes. Indeed, action languages provide a limited support for process definition, in terms of workflow constructs, and they lack a clear mapping from standard (ontology and process) modeling languages.

The main objective of this paper is to design a framework for representing and reasoning about business process knowledge from both the procedural and ontological point of views. To achieve this goal, we do *not* propose yet another business process modeling language, but we provide a framework based on Logic Programming (LP) [38] for reasoning about process-related knowledge expressed by means of de-facto standards for BP modeling, like BPMN [46], and ontology definition, like OWL [43]. We define a rule-based procedural semantics for a relevant fragment of BPMN, by following an approach inspired by the *Fluent Calculus* [61], and we extend it in order to take into account OWL annotations that describe preconditions and effects of activities and events occurring within a BP. In particular, we integrate our procedural BP semantics with the OWL 2 RL profile thanks to a common grounding in LP. OWL 2 RL is indeed a fragment of the OWL ontology language that has a suitable rule-based presentation, thus constituting an excellent compromise between expressivity and efficiency.

The contributions of this paper can be summarized as follows.

After presenting the preliminaries in Section 2, we propose, in Section 3, a revised and extended version of the Business Process Abstract Language (BPAL) [16, 56], a process ontology for modeling the procedural semantics of a BP regarded as a workflow. To this end we introduce an axiomatization to cope with a relevant fragment of the BPMN 2.0 specification, allowing us to deal with a large class of process models.

We then propose, in Section 4, an approach for the semantic annotation of BP models, where BP elements are described by using an OWL 2 RL ontology.

In Section 5 we provide a general verification mechanism by integrating the temporal logic CTL [15] within our framework, in order to analyze properties of the *states* that the system can reach, by taking into account both the control-flow and the semantic annotation.

In Section 6 we show how a repository of semantically enriched BPs can be organized in a Business Process Knowledge Base (BPKB), which, due to the common representation of its components in LP, provides a uniform and formal framework that enables logical inference. We then discuss how, by using state-of-the-art LP systems, we can perform some very sophisticated reasoning tasks, such as verification, querying and trace compliance checking, that combine both the procedural and the domain knowledge relative to a BP.

In Section 7 we provide the computational characterization of the reasoning services that can be performed on top of a BPKB, showing in particular that, for a large class of them, advanced resolution strategies (such as SLG-Resolution [14]) guarantee an efficient, sound and complete procedure.

In Section 8 we describe the implemented tool, which provides a graphical user interface to support the semantic BP design, and a reasoner, developed in XSB Prolog [58], able to operate on the BPKB. We also report on an evaluation of the system performance, demonstrating that complex reasoning tasks can be performed on business process of small-to-medium size in an acceptable amount of time and memory resources.

In Section 9 we compare our work to related approaches and in the concluding section we give a critical discussion of our approach, along with directions for future work.

2 Preliminaries

In order to clarify the terminology and the notation used throughout this paper, in this section we recall some basic notions related to the BPMN notation [46], Description Logics [4] as well as foundations of the OWL [43] standard, and Logic Programming [38].

2.1 BPMN

Business Process Modeling and Notation (BPMN) [46] is a graphical language for BP modeling, standardized by the OMG (`http://www.omg.org`). The primary goal of BPMN is to provide a standard notation readily understandable by all business stakeholders, which include the business analysts who create and refine the processes, the technical developers responsible for their implementation, and the business managers who monitor and manage them.

A BPMN model is defined through a Business Process Diagram (BPD), which is a kind of flowchart incorporating constructs to represents the control flow, data flow, resource allocation (i.e., how the work is assigned to the participants), and exception handling (i.e., how erroneous behavior can be handled and compensated). We will briefly overview the core BPMN constructs referring to the example in Figure 1.

The constructs of BPMN are classified as flow objects, artifacts, connecting objects, and swimlanes.

Flow objects are partitioned into *activities* (represented as rounded rectangles), *events* (represented as circles), and *gateways* (represented as diamonds). Activities are a generic way of representing some kind of work performed within the process, and can be *tasks* (i.e.,

atomic activities such as *create_order*) or *compound activities* corresponding to the execution of entire sub-processes (e.g., *create_order*). Events denote something that "happens" during the enactment of a business process, and are classified as *start events*, *intermediate events*, and *end events* which can start (e.g., *s*), suspend (e.g., *ex*), or end (e.g., *e*) the process enactment. An intermediate event, such as *ex*, attached to the boundary of an activity models exception handling. Gateways model the branching and merging of the control flow. There are several types of gateways in BPMN, each of which may be used as a *branch* gateway if it has multiple outgoing flows, or a *merge* gateway if it has multiple incoming flows. The split and join behavior depends on the semantics associated to each type of gateway. Exclusive branch gateways (e.g., *g1*) are decision points where exactly one of a set of mutually exclusive alternative flows is selected, while an exclusive merge gateway (e.g., *g2*) merges two incoming flows into a single one. Parallel branch gateways (e.g., *g7*) create parallel threads of execution, while parallel merge gateways (e.g., *g8*) synchronize concurrent flows. Inclusive branch gateways (e.g., *g3*) are decision points where at least one of a set of non-exclusive alternative flows is selected, while an inclusive merge gateway (e.g., *g4*) is supposed to be able to synchronize a varying number of threads, i.e., it is executed only when at least one of its predecessors has been executed and no other will be eventually executed[1].

Connecting objects are *sequence flows* (e.g., the directed edge between *g1* and *g3*) and *associations* (e.g., the dashed edge between *create_order* and *order*). A sequence flow links two flow objects and denotes a control flow relation, i.e., it states that the control flow can pass from the source to the target object. An *association* is used to associate *artifacts* (i.e., data objects) with flow objects, and its direction defines if a data object is used as an input (e.g., *order* is an input of *accept_order*) or it is an output (e.g., *order* is an output of *create_order*) of some flow element.

Swimlanes are used to model participants, i.e., a generic notion representing a role within a company (e.g., Sales Clerk), a department (e.g., Finance) or a business partner (e.g., Courier), which is assigned to the execution of a collection of activities.

2.2 Description Logics and Rule-based OWL Ontologies

Description Logics (DLs) [4] are a family of knowledge representation languages that can be used to represent the knowledge of an application domain in a structured and formally well-understood way. DLs are typically adopted for the definition of ontologies since on the one hand, the important notions of the domain are described by *concept descriptions*, i.e., expressions that are built from atomic concepts (usually thought as sets of individuals, e.g., $Person$) and atomic roles (relations between concepts, e.g., $worksFor$) using the concept and role constructors provided by the particular DL (e.g., $Person \sqcap \exists worksFor.Company$, that is, the set of *persons* who work for a *company*). On the other hand, DLs correspond to decidable fragments of classical first-order logic (FOL), and thus are equipped with a formal, logic-based semantics that makes such languages suitable for automated reasoning.

Typically, Description Logics are used for representing a TBox (terminological box) and the ABox (Assertional Box). The TBox describes concept (and role) hierarchies, (e.g., $Employee \sqsubseteq Person \sqcap \exists worksFor.Company$), while the ABox contains assertions about individuals (e.g., $john : Employee$).

[1]For sake of completeness, BPMN provides two more types of gateways, which we do not exemplify, namely, the event-based and the complex gateway.

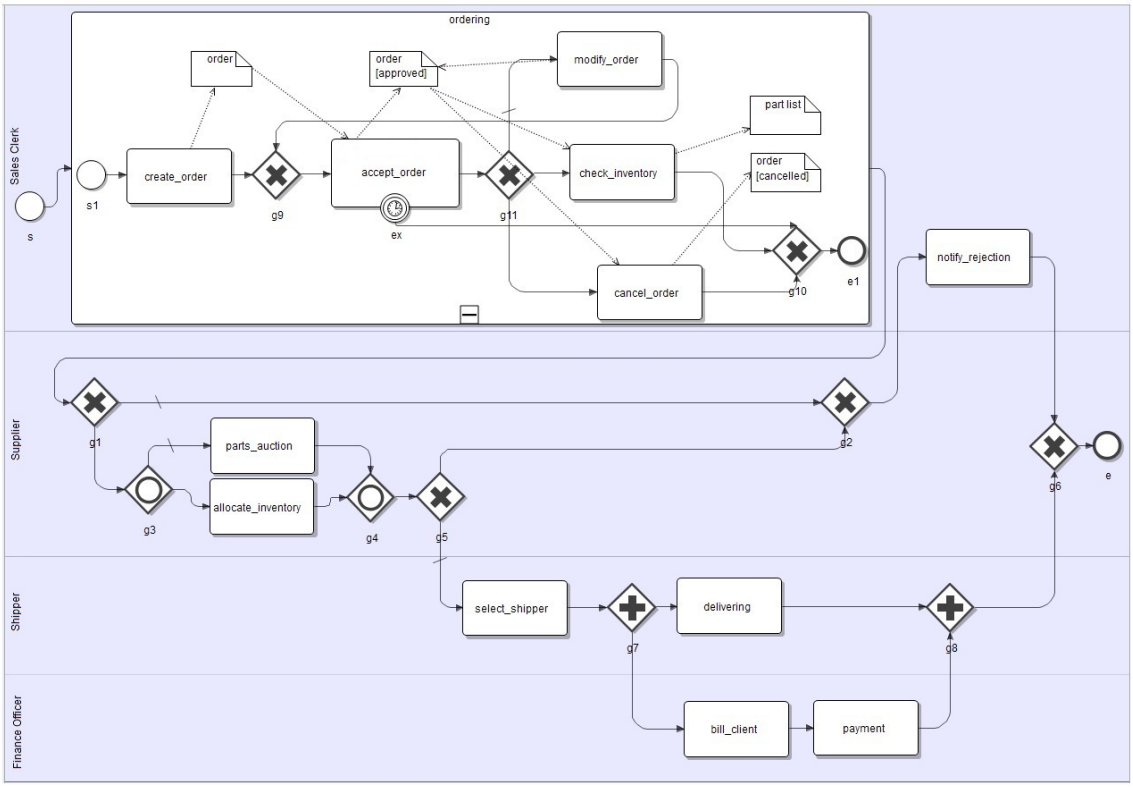

Figure 1: Handle Order Business Process

The growing interest in the Semantic Web vision [7], where Knowledge Representation techniques are adopted to make resources machine-interpretable by "intelligent agents", has pushed the standardization of languages for ontology and meta-data sharing over the (semantic) web. Among these, one of the most promising standards is the Ontology Web Language (OWL) [43], formally grounded in DLs, proposed by the Web Ontology Working Group of W3C. OWL is *syntactically* layered on RDF [34] and RDFS [10], and can be considered as an extension of RDFS in terms of modeling capabilities and reasoning facilities. The underlying data model (derived from RDF) is based on statements (or RDF triples) of the form $< subject, property, object >$, which allow us to describe a resource (*subject*) in terms of named relations (*properties*). Values of named relations (i.e. objects) can be URIrefs of Web resources or literals, i.e., representations of data values (such as integers and strings).

Table 1 shows, for some OWL statements, the corresponding DL notations and FOL formulae, where C and D are concepts (OWL classes), P and Q are roles (OWL properties), a and b are constants, and x and y are variables.

The recent OWL 2 specification defines profiles that correspond to syntactic subsets of OWL, each of which is designed to trade some expressive power for efficiency of reasoning. In particular, we consider OWL 2 RL, closely related to the Horn fragment of FOL, which is based on *Description Logic Programs* [28] and *pD** [59]. The use of OWL 2 RL allows us to take advantage of the efficient resolution strategies developed for logic programs, in order to perform the reasoning tasks typically supported by Description Logics reasoning systems, such as concept subsumption and ontology consistency. Indeed, the semantics of OWL 2 RL is defined through a partial axiomatization of the OWL 2 RDF-Based

Table 1: Main OWL statements and FOL equivalence

OWL Axiom	***DL Expression***	***FOL Formula***
a type C	$a : C$	$C(a)$
a P b	$(a, b) : P$	$P(a, b)$
C subClassOf D	$C \sqsubseteq D$	$\forall x.C(x) \rightarrow D(x)$
C disjointWith D	$C \sqsubseteq \neg D$	$\forall x.C(x) \rightarrow \neg D(x)$
P domain C	$\top \sqsubseteq \forall P^{-}.C$	$\forall x,y.P(x, y) \rightarrow C(x)$
P range C	$\top \sqsubseteq \forall P.C$	$\forall x,y.P(x, y) \rightarrow C(y)$
transitiveProperty P	$P^{+} \sqsubseteq P$	$\forall x,y,z.(P(x, y) \wedge P(y, z)) \rightarrow P(x, z)$
functionalProperty P	$\top \sqsubseteq \leq 1\ P$	$\forall x,y,z.(P(x, y) \wedge P(x, z)) \rightarrow y = z$
P inverseOf Q	$P \equiv Q^{-}$	$\forall x,y.P(x, y) \leftrightarrow Q(y, x)$
OWL Constructor		
C intersectionOf D	$C \sqcap D$	$C(x) \wedge D(x)$
C unionOf D	$C \sqcup D$	$C(x) \vee D(x)$
P allValuesFrom C	$\forall P.C$	$\forall y.P(x, y) \rightarrow C(y)$
P someValuesFrom C	$\exists P.C$	$\exists y.P(x, y) \wedge C(y)$
complementOf D	$\neg D$	$\neg D(x)$

Semantics in the form of first-order implications (OWL 2 RL/RDF rules), and constitutes an upward-compatible extension of RDF and RDFS.

OWL 2 RL ontologies are modeled by means of the ternary predicate $t(s, p, o)$ representing an OWL statement with subject s, predicate p and object o. For instance, the assertion $t(a, \textit{rdfs:subClassOf}, b)$ represents the inclusion axiom $a \sqsubseteq b$. Reasoning on triples is supported by OWL 2 RL/RDF rules of the form $t(s, p, o) \leftarrow t(s_1, p_1, o_1) \wedge \cdots \wedge t(s_n, p_n, o_n)$. Table 2 shows some of the rules of the OWL 2 RL/RDF rule-set. According to the terminology we will introduce in the next section, this rule set is a *definite logic program*.

Table 2: Excerpt of the OWL 2 RL/RDF rule-set

Transitive subsumption	$t(C_1, \textit{rdfs:subClassOf}, C_3) \leftarrow t(C_1, \textit{rdfs:subClassOf}, C_2) \wedge$ $t(C_2, \textit{rdfs:subClassOf}, C_3)$
Inheritance	$t(X, \textit{rdf:type}, C_2) \leftarrow t(C_1, \textit{rdfs:subClassOf}, C_2) \wedge t(X, \textit{rdf:type}, C_1)$ $t(X, \textit{rdf:type}, C_2) \leftarrow t(C_1, \textit{owl:equivalentClass}, C_2) \wedge t(X, \textit{rdf:type}, C_2)$
Domain	$t(X, \textit{rdf:type}, C) \leftarrow t(P, \textit{rdfs:domain}, C) \wedge t(X, P, O)$
Range	$t(Y, \textit{rdf:type}, C) \leftarrow t(P, \textit{rdfs:range}, C) \wedge t(S, P, Y)$
Transitivity	$t(X, P, Z) \leftarrow t(P, \textit{rdf:type}, \textit{owl:TransitiveProperty}) \wedge t(X, P, Y) \wedge t(Y, P, Z)$
Subsumption of existential formulae	$t(C_1, \textit{rdfs:subClassOf}, C_2) \leftarrow t(C_1, \textit{owl:someValuesFrom}, D_1) \wedge$ $t(C_1, \textit{owl:onProperty}, P) \wedge t(C_2, \textit{owl:someValuesFrom}, D_2) \wedge$ $t(C_2, \textit{owl:onProperty}, P) \wedge t(D_1, \textit{rdfs:subClassOf}, D_2)$
Intersection	$t(C, \textit{rdfs:subClassOf}, D) \leftarrow t(C, \textit{owl:intersectionOf}, I) \wedge D \in I$
Disjointness	$\bot \leftarrow t(C_1, \textit{owl:disjointWith}, C_2) \wedge t(X, \textit{rdf:type}, C_1) \wedge t(X, \textit{rdf:type}, C_2)$

2.3 Logic programming

We briefly recall the basic notions of Logic Programming. In particular, we will consider the class of *locally stratified logic programs*, or *stratified programs*, for short, and their standard semantics defined by the *perfect model*. (Recall that all major declarative semantics of logic programs coincide on stratified programs.) This class of logic programs is expres-

sive enough to represent several complementary pieces of knowledge related to business processes, such as the syntactic structure of the control flow, the operational semantics, the ontology-based properties, and the temporal properties of the execution. For more details about LP we refer to [38, 2].

A *term* is either a *constant*, or a *variable*, or an expression of the form $f(t_1, \ldots, t_m)$, where f is a function symbol and $t_1, \ldots, t_m$ are terms. An *atom* is a formula of the form $p(t_1, \ldots, t_m)$, where p is a *predicate symbol* and $t_1, \ldots, t_m$ are terms. A *literal* is either an atom or a negated atom. A *rule* is a formula of the form $A \leftarrow L_1 \wedge \ldots \wedge L_n$, where A is an atom (the *head* of the rule) and $L_1 \wedge \ldots \wedge L_n$ is a conjunction of literals (the *body* of the rule). If $n = 0$ we call the rule a *fact*. A rule (term, atom, literal) is *ground* if no variables occur in it. A *logic program* is a set of rules. A *definite* program is a logic program with no negated atoms in the body of its rules. For a logic program P, by *ground*(P) we denote the set of ground instances of rules in P.

Let B_P denote the *Herbrand base* for P, that is, the set of ground atoms that can be constructed in the language of program P. An (Herbrand) *interpretation* I is a subset of B_P. A ground atom A is true in I if $A \in I$. A ground negated atom $\neg A$ is true in I if $A \notin I$. A ground rule $A \leftarrow L_1 \wedge \ldots \wedge L_n$ is true in I if either A is true in I or, for some $i \in \{1, \ldots, n\}$, L_i is not true in I. An interpretation is a *model* of P if all rules in *ground*(P) are true in I. Every definite program has a *least* Herbrand model. However, this property does not hold for general logic programs.

A *(local) stratification* is a function σ from the Herbrand base B_P to the set of all countable ordinals [2, 50]. However, for the purposes of this paper it will be enough to consider stratification functions from B_P to the set $\mathbb{N}$ of the natural numbers. For a ground atom A, $\sigma(A)$ is called the *stratum* of A. A stratification σ extends to negated atoms by taking $\sigma(\neg A) = \sigma(A) + 1$. A ground rule $A \leftarrow L_1 \wedge \ldots \wedge L_n$ is *stratified with respect to* σ if, for $i = 1, \ldots, n$, $\sigma(A) \geq \sigma(L_i)$. A program P is stratified with respect to σ if every rule in *ground*(P) is. Finally, a logic program is *stratified* if it is stratified with respect to some stratification function.

The perfect model of P, denoted *Perf*(P), is defined as follows. Let P be stratified with respect to σ. For every $n \in \mathbb{N}$, let S_n be the set of rules in *ground*(P) whose head has stratum n. Thus, *ground*$(P) = \bigcup_{n \in \mathbb{N}} S_n$. We define a sequence of interpretations as follows: (i) M_0 is the least model of S_0 (note that S_0 is a definite program), and (ii) M_{n+1} is the least model of S_n that contains M_n. The perfect model of P, is defined as *Perf*$(P) = \bigcup_{n \in \mathbb{N}} M_n$. (Here we are using the simplifying assumption that the codomain of the stratification function is $\mathbb{N}$.)

The operational semantics of logic programs is based on the notion of derivation, which is constructed by *SLD-resolution* augmented with the *Negation as Failure* rule [38]. Given a stratified program P, we will define below the *one-step derivation* relation $Q_1 \xrightarrow{\theta} Q_2$, where Q_1, Q_2 are *queries*, that is, conjunctions of literals, and θ is a substitution. The definition of one-step derivation relation depends on the following notions. A *derivation* for a query Q_0 with respect to P is a sequence $Q_0 \xrightarrow{\theta_1} \ldots \xrightarrow{\theta_n} Q_n$ $(n \geq 1)$. We will omit the reference to P when clear from the context. A derivation is *successful* if its last query is the empty conjunction *true*. A query *succeeds* if there exists a successful derivation for it. A query *fails* if it does not succeed. The one-step derivation relation is defined by the following two derivation rules.

(P) Let $A \wedge Q$ be a query, where A is an atom. Suppose that $H \leftarrow K_1 \wedge \ldots \wedge K_m$ $(m \geq 0)$ is a rule in P such that A is unifiable with H via a most general unifier θ [38]. Then $A \wedge Q \xrightarrow{\theta} (K_1 \wedge \ldots \wedge K_m \wedge Q)\theta$.

(N) Let $\neg A \wedge Q$ be a query, where A is a ground atom. Suppose that A fails. Then $\neg A \wedge Q \xrightarrow{\epsilon} Q$, where ϵ is the identity substitution.

Note that in the definition of a derivation we assume the left-to-right selection rule for literals. Note also that, in rule (N) the one-step derivation from $\neg A \wedge Q$ refers to the set of all derivations from A (to show that A fails). However, this definition is well-founded because the program P is stratified. We say that a query Q is *generable* from a query Q_0 if there exists either a derivation $Q_0 \xrightarrow{\theta_1} \ldots \xrightarrow{\theta_n} Q$ or a derivation $Q_0 \xrightarrow{\theta_1} \ldots \xrightarrow{\theta_n} \neg A \wedge Q_n$ and Q is generable from A. An *answer* for a query Q_0 is a substitution θ such that there exists a successful derivation $Q_0 \xrightarrow{\theta_1} \ldots \xrightarrow{\theta_n} \mathit{true}$ and θ is the restriction of the composition $\theta_1 \ldots \theta_n$ to the variables occurring in Q_0. A query Q_0 *flounders* if there exists a query Q generable from Q_0 such that the leftmost literal of Q is a non ground negated atom.

The operational semantics is *sound* and *complete* with respect to the perfect model semantics for queries that do not flounder. Indeed, it can be shown that (see, for instance, [50, 2]), given a program P and an atom A_0 that does not flounder with respect to P, then: (1) if A_0 succeeds with answer θ, then every ground instance of $A_0\theta$ belongs to *Perf*(P), and (2) if $A_0\theta$ belongs to *Perf*(P) for some substitution θ, then A_0 succeeds with an answer which is more general than θ.

The definition of a derivation given above is quite abstract and not fully constructive. In particular, the application of rule (N) requires to test that an atom has no successful derivations, and this property is undecidable in the general case. Thus, an effective query evaluation strategy depends on the concrete way derivations are constructed.

A well-known difficulty of the evaluation strategy based on depth-first search is that infinite derivations may be constructed, even in cases where a finite set of atoms (modulo variants) is derived from a given initial query. In particular, this nonterminating behavior can occur for stratified *Datalog* programs, that is, function free stratified programs.

In order to avoid this difficulty, in this paper we adopt *SLG-resolution*, a query evaluation mechanism that implements SLD resolution with Negation as Failure by means of *tabling* [14]. During the construction of the derivations for a given atom A_0, a *table* is maintained to record the answers to A_0 and to the atoms generated from A_0. The tabled answers are used the next time an atom is generated, and hence no atom is evaluated more than once. Thus, SLG-resolution is able to compute in finite time all answers to a query, if a finite set of atoms is generated and a finite set of answers for those atoms exists. In particular, SLG-resolution always terminates and is able to compute all answers for queries to stratified Datalog programs.

3 Rule-based Representation of BP Schemas

In this section we introduce a formal representation of business processes by means of the notion of *Business Process Schema* (BPS). A BPS, its meta-model, and its procedural (or *behavioral*) semantics will all be specified by sets of rules, for which we adopt the standard notation and semantics of LP (see Section 2.3).

3.1 Introducing BPAL

The Business Process Abstract Language (BPAL) introduces a language conceived to provide a declarative modeling method capable of fully capturing procedural knowledge in a business process. BPAL constructs are common to the most used and widely accepted BP modeling languages (e.g., BPMN [46], UML activity diagrams [47], EPC [33]) and, in particular, it is based on the BPMN 2.0 specification [46].

Formally, a (set of) BPS(s) $\mathcal{B}$ is specified by a set of ground facts of the form $p(c_1, \ldots, c_n)$, where $c_1, \ldots, c_n$ are constants denoting flow elements (e.g., activities, events, and gateways) and p is a predicate symbol. In Table 3 we list some of the BPAL predicates, and in Table 4 we exemplify their usage reporting the translation of the *Handle Order* process (*ho* for short) depicted in Figure 1 as a BPAL BPS. An extended discussion can be found in [16, 55].

Table 3: Excerpt of the BPAL language

Construct	**Description**
bp(p,s,e)	p is a process, with entry-point s and exit-point e
element(x)	x is a flow *element* occurring in some process
relation(x,y,p)	the *elements* x and y are in relation in the process p
task(a)	a is an atomic activity
event(e)	e is an event
exception(e,a,p)	the intermediate event e (an exception) is attached to the activity a
comp_act(a,s,e)	a is a compound activity, with entry-point s and exit-point e
seq(el1,el2,p)	a sequence flow relation is defined between *el1* and *el2* in p
par_branch(g)	the execution of g enables all the successor flow elements
par_merge(g)	g waits for the completion of all the predecessor flow elements
exc_branch(g)	the execution of g enables one of the successor flow elements
exc_merge(g)	g waits for the completion of one of the predecessor flow elements
inc_branch(g)	the execution of g enables at least one of its successors
inc_merge(g)	g waits for the completion of the predecessor flow elements that will be eventually executed
item(i)	i is a data element
input(a,i,p)	the activity a uses as input the data element i in the process p
output(a,i,p)	the activity a uses as output the data element i in the process p
participant(part)	*part* is a participant
assigned(a,part,p)	the activity a is assigned to the participant *part* in the process p

Table 4: BPS representing the Handle Order process

bp(ho,s,e)	*seq(s, ordering, ho)*	*comp_act(ordering, s_1, e_1)*
seq(ordering,g1,ho)	*seq(g1,g2, ho)*	*assigned(ordering,sales_clerk,ho)*
seq(g1,g3,ho)	*seq(g3,parts_auction,ho)*	*assigned (delivering,shipper,ho)*
seq(g3,allocate_inventory,ho)	*seq(parts_auction,g4,ho)*	*seq(s1,create_order,ordering)*
seq(allocate_inventory,g4,ho)	*seq(g4,g5,ho)*	*seq(create_order,g9,ordering)*
seq(g5,g2,ho)	*seq(g5,select_shipper,ho)*	*seq(g9,accept_order,ordering)*
seq(g2,notify_rejection,ho)	*seq(select_shipper,g7,ho)*	*exception(ex,accept_order,ordering)*
seq(notify_rejection,g6,ho)	*seq(g6,e,ho)*	*seq(ex,g10,ordering)*
exc_branch(g1)	*participant(sales_clerk)*	*input(accept_order,order,ordering)*
inc_branch(g3)	*task(create_order)*	*output(create_order,order,ordering)*
par_branch(g7)	*item(order)*	...

Our formalization also includes in $\mathcal{B}$ a set of rules that represents the *meta-model*, defining *i*) hierarchical relationships among the BPAL predicates, e.g., $activity(x) \leftarrow task(x)$; *ii*) disjointness relationships among BPAL elements, e.g., $\bot \leftarrow activity(x) \wedge event(x)$; *iii*) structural properties which regard a BPS as a directed graph, where edges correspond to sequence and item flow relations. A first set of structural properties represents constraints that should be verified by a *well-formed* BPS, i.e., syntactically correct BPS: (1) every process is assigned to a unique start event and to a unique end event; (2) every flow element occurs on a path from the start event to the end event; (3) start events have no predecessors and end events have no successors; (4) branch gateways have exactly one predecessor and at least two successors, while merge gateways have at least two predecessors and exactly one successor; (5) activities and intermediate events have exactly one predecessor and one successor; (6) there are no cycles in the hierarchy of compound activities.

Finally, other meta-model properties are related to the notions of path and reachability between flow elements, such as the following ones, which will be used in the sequel: $seq^{+}(E_1, E_2, P)$, representing the transitive closure of the *sequence flow* relation, and $n_reachable(E_1, E_2, E_3, P)$, which holds if there is a path in P between E_1 and E_2 not including E_3, i.e.:

$n_reachable(X, Y, N, P) \leftarrow seq(X, Y, P) \wedge \neg Y = N$
$n_reachable(X, Y, N, P) \leftarrow seq(X, Z, P) \wedge \neg Z = N \wedge n_reachable(Z, Y, N, P)$

With respect to the framework introduced in [16, 55], here we consider unstructured cyclic workflows whose behavioral semantics will be introduced in the following.

3.2 Behavioral Semantics

Now we present a formal definition of the behavioral semantics, or *enactment*, of a BPS, by following an approach inspired by the *Fluent Calculus*, a well-known calculus for action and change (see [61] for an introduction), which is formalized in Logic Programming.

In the Fluent Calculus, the state of the world is represented as a collection of *fluents*, i.e., terms representing atomic properties that hold at a given instant of time. An action, also represented as a term, may cause a change of state, i.e., an update of the collection of fluents associated with it. Finally, a *plan* is a sequence of actions that leads from the initial to the final state. For states we use set notation (here we depart from [61], where an associative-commutative operator is used for representing collections of fluents). A fluent is an expression of the form $f(a_1, \ldots, a_n)$, where f is a fluent symbol and $a_1, \ldots, a_n$ are constants or variables. In order to model the behavior of a BPS, we represent states as *finite sets* of ground fluents. We take a closed-world interpretation of states, that is, we assume that a fluent F, different from *true*, holds in a state S iff $F \in S$. Our set-based representation of states relies on the assumption that the BPS is *safe*, that is, during its enactment there are no concurrent executions of the same flow element [64]. This assumption enforces that the set of states reachable by a given BPS is finite. A *fluent expression* is built inductively from fluents, the binary function symbol *and*, and the unary function symbol *not*. The satisfaction relation assigns a truth value to a fluent expression with respect to a state. This relation is encoded by a predicate $holds(F, S)$, which holds if the fluent expression F is true in the state S. We also introduce a constant symbol *true*, such that $holds(true, S)$ holds for every state S. Accordingly to the closed-world interpretation given to states, the satisfaction relation is defined by the following rules:

$holds(F, S) \leftarrow F = true$
$holds(F, S) \leftarrow F \in S$
$holds(not(F), S) \leftarrow \neg holds(F, S)$
$holds(and(F_1, F_2), S) \leftarrow holds(F_1, S) \wedge holds(F_2, S)$

Note that, by the perfect model semantics, reflecting the closed-world assumption, for any fluent F different from *true*, $not(F)$ holds in a state S iff $F \notin S$.

We will consider the following two kinds of fluents:

- $cf(E_1, E_2, P)$, which means that the flow element E_1 has been executed and the successor flow element E_2 is waiting for execution, during the enactment of the process P (*cf* stands for *control flow*);
- $en(A, P)$, which means that the activity A is being executed during the enactment of the process P (*en* stands for *enacting*).

To clarify our terminology note that, when a flow element E_2 is waiting for execution, E_2 might not be enabled to execute, because other conditions need to be fulfilled, such as those depending on the synchronization with other flow elements (see, in particular, the semantics of merging behaviors below).

We assume that the execution of an activity has a beginning and a completion (although we do not associate a *duration* with activity execution), while the other flow elements execute instantaneously. Thus, we will consider two kinds of actions: $begin(A)$ which starts the execution of an activity A, and $complete(E)$, which represents the completion of the execution of a flow element E (possibly, an activity). The change of state determined by the execution of an action will be formalized by a relation $result(S_1, A, S_2)$, which holds if action A can be executed in state S_1 leading to state S_2. For defining the relation $result(S_1, A, S_2)$ the following auxiliary predicates will be used: (i) $update(S_1, T, U, S_2)$, which holds if $S_2 = (S_1 - T) \cup U$, where S_1, T, U, and S_2 are sets of fluents, and (ii) $setof(F, C, S)$, which holds if S is the set of ground instances of fluent F such that condition C holds.

The relation $r(S_1, S_2)$ holds if a state S_2 is *immediately reachable* from a state S_1, that is, some action A can be executed in state S_1 leading to state S_2:

$r(S_1, S_2) \leftarrow result(S_1, A, S_2)$

We say that a state S_2 is *reachable* from a state S_1 if there is a finite, possibly empty, sequence of actions from S_1 to S_2, that is, $reachable_state(S_1, S_2)$ holds, where the relation *reachable_state* is is the reflexive-transitive closure of r.

In the rest of this section we present a fluent-based formalization of the behavioral semantics of a BPS as a set of rules $\mathcal{T}$, partially reported in Table 5. The proposed formal semantics is focused on a core of the BPMN language and it mainly refers to its semantics, as described (informally) in the most recent specification of the language [46]. Most of the constructs considered here (e.g., parallel or exclusive branching/merging) have the same interpretation in most workflow languages. However, when different interpretations are given, e.g., in the case of inclusive merge, we stick to the BPMN one.

3.2.1 Activity and Event Execution

The enactment of a process P begins with the execution of the associated start event E in a state where the fluent $cf(start, E, P)$ holds, being *start* a reserved constant. After the execution of the start event, its unique successor waits for execution (Rule E1). The

execution of an end event leads to the final state of a process execution, in which the fluent $cf(E, end, P)$ holds, where E is the end event associated with the process P and *end* is a reserved constant (Rule E2).

According to the informal semantics of BPMN, *intermediate events* are intended as instantaneous patterns of behavior that are registered at a given time point. Thus, we formally model the execution of an intermediate event as a single state transition, as defined in Rule E3. Intermediate events in BPMN can also be attached to activity boundaries to model exceptional flows. Upon occurrence of an *exception*, the execution of the activity is interrupted, and the control flow moves along the sequence flow that leaves the event (Rule E4).

The execution of an activity is enabled to begin after the completion of its unique predecessor flow element. The effects of the execution of an activity vary depending on its type (i.e., atomic task or compound activity). The beginning of an atomic task A is modeled by adding the $en(A, P)$ fluent to the state (Rule A1). At the completion of A, the $en(A, P)$ fluent is removed and the control flow moves on to the unique successor of A (Rule A2). The execution of a compound activity, whose internal structure is defined as a process itself, begins by enabling the execution of the associated *start event* (Rule A3), and completes after the execution of the associated *end event* (Rule A4).

3.2.2 Branching Behaviors

When a branch gateway is executed, a subset of its successors is selected for execution. We consider here exclusive, inclusive, and parallel branch gateways.

An exclusive branch leads to the execution of exactly one successor (Rule B1), while an inclusive branch leads to the concurrent execution of a non-empty subset of its successors (Rule B2). The set of successors of exclusive or inclusive decision points may depend on *guards*, i.e., conditions that usually take the form of tests on the value of the items that are passed between the activities. While Rules B1-B2 formalize a nondeterministic choice among the successors of a decision point, in Section 4.3 guard expressions will be included in the framework in the form of fluent expressions whose truth value is tested with respect to the current state. Finally, a parallel branch leads to the concurrent execution of all its successors (Rule B3).

3.2.3 Merging Behaviors

An exclusive merge can be executed whenever at least one of its predecessors has been executed (Rule X1).

For the inclusive merge several operational semantics have been proposed, due to the complexity of its non-local semantics (see e.g., [33, 65]). An inclusive merge is supposed to be able to synchronize a varying number of threads, i.e., it is executed only when $n(\geq 1)$ predecessors have been executed and no other will be eventually executed. Here we refer to the semantics described in [65] adopted by BPMN, stating that (Rule O1) an inclusive merge M can be executed if the following two conditions hold (Rules O2, O3):

(1) at least one of its predecessors has been executed,

(2) for each non-executed predecessor X, there is no flow element U which is waiting for execution and is *upstream* X. The notion of being upstream captures the fact that U may lead to the execution of X, and is defined as follows. A flow element U is

upstream X if (Rules O4, O5): *a)* there is a path from U to X not including M, and *b)* there is no path from U to an executed predecessor of M not including M.

Finally, a parallel merge can be executed if all its predecessors have been executed as defined in Rule P1, where $exists_non_executed_pred(M, P, S_1)$ holds if there exists no predecessor of M which has not been executed in state S_1 (Rule P2).

3.2.4 Item Flow

BP modeling must be able to represent the physical and the information items that are produced and consumed by the various activities during the execution of a process. For the formalization of the *item flow* semantics, we commit to the BPMN standard, where the so-called *data objects* are used to store information created and read by activities.

In our approach items are essentially regarded as variables, and hence there is a single instance of a given item any time during the execution that may be (over-)written by some activity. We consider two main types of relations between activities and items. First of all, an activity may *use* a particular item (*input* relation). This implies that the item is

Table 5: Fragment of the behavioral semantics of the BPAL language

(E1) $result(S_1, complete(E), S_2) \leftarrow start_event(E) \wedge holds(cf(start, E, P), S_1) \wedge seq(E, X, P) \wedge update(S_1, \{cf(start, E, P)\}, \{cf(E, X, P)\}, S_2)$

(E2) $result(S_1, complete(E), S_2) \leftarrow end_event(E) \wedge holds(cf(X, E, P), S_1) \wedge update(S_1, \{cf(X, E, P)\}, \{cf(E, end, P)\}, S_2)$

(E3) $result(S_1, complete(E), S_2) \leftarrow int_event(E) \wedge holds(cf(X, E, P), S_1) \wedge seq(E, Y, P) \wedge update(S_1, \{cf(X, E, P)\}, \{cf(E, Y, P)\}, S_2)$

(E4) $result(S_1, complete(E), S_2) \leftarrow exception(E, A, P) \wedge int_event(E) \wedge holds(en(A, P), S_1) \wedge seq(E, Y, P) \wedge update(S_1, \{en(A, P)\}, \{cf(E, Y, P)\}, S_2)$

(A1) $result(S_1, begin(A), S_2) \leftarrow task(A) \wedge holds(cf(X, A, P), S_1) \wedge update(S_1, \{cf(X, A, P)\}, \{en(A, P)\}, S_2)$

(A2) $result(S_1, complete(A), S_2) \leftarrow task(A) \wedge holds(en(A, P), S_1) \wedge seq(A, Y, P) \wedge update(S_1, \{en(A, P)\}, \{cf(A, Y, P)\}, S_2)$

(A3) $result(S_1, begin(A), S_2) \leftarrow comp_act(A, S, E) \wedge holds(and(cf(X, A, P), not(en(A, P))), S_1) \wedge update(S_1, \{cf(X, A, P)\}, \{cf(start, S, A), en(A, P)\}, S_2)$

(A4) $result(S_1, complete(A), S_2) \leftarrow comp_act(A, S, E) \wedge holds(and(cf(E, end, A), en(A, P)), S_1) \wedge seq(A, Y, P) \wedge update(S_1, \{en(A, P), cf(E, end, A)\}, \{cf(A, Y, P)\}, S_2)$

(B1) $result(S_1, complete(B), S_2) \leftarrow exc_branch(B) \wedge holds(cf(X, B, P), S_1) \wedge seq(B, Y, P) \wedge update(S_1, \{cf(X, B, P)\}, \{cf(B, Y, P)\}, S_2)$

(B2) $result(S_1, complete(B), S_2) \leftarrow inc_branch(B) \wedge holds(cf(X, B, P), S_1) \wedge setof(cf(B, Y, P), seq(B, Y, P), Succ) \wedge subseteq(SubSucc, Succ) \wedge \neg emptyset(SubSucc) \wedge update(I, \{cf(X, B, P)\}, SubSucc, S_2)$

(B3) $result(S_1, complete(B), S_2) \leftarrow par_branch(B) \wedge holds(cf(X, B, P), S_1) \wedge setof(cf(B, Y, P), seq(B, Y, P), Succ) \wedge update(S_1, \{cf(X, B, P)\}, Succ, S_2)$

(X1) $result(S_1, complete(M), S_2) \leftarrow exc_merge(M) \wedge holds(cf(A, M, P), S_1) \wedge seq(M, Y, P) \wedge update(S_1, \{cf(A, M, P)\}, \{cf(M, Y, P)\}, S_2)$

(O1) $result(S_1, complete(M), S_2) \leftarrow inc_merge(M) \wedge enabled_im(M, S_1, P) \wedge seq(M, Y, P) \wedge setof(cf(X, M, P), holds(cf(X, M, P), S_1), PredM) \wedge update(S_1, PredM, \{cf(M, Y, P)\}, S_2)$

(O2) $enabled_im(M, S_1, P) \leftarrow holds(cf(X, M, P), S_1) \wedge \neg exists_upstream(M, S_1, P)$

(O3) $exists_upstream(M, S_1, P) \leftarrow seq(X, M, P) \wedge holds(not(cf(X, M, P)), S_1) \wedge holds(cf(Y, U, P), S_1) \wedge upstream(U, X, M, S_1, P)$

(O4) $upstream(U, X, M, S_1, P) \leftarrow n_reachable(U, X, M, P) \wedge \neg exists_path(U, M, S_1, P)$

(O5) $exists_path(U, M, S_1, P) \leftarrow holds(cf(K, M, P), S_1) \wedge n_reachable(U, K, M, P)$

(P1) $result(S_1, complete(M), S_2) par_merge(M) \wedge \neg exists_non_executed_pred(M, P, S_1) \wedge seq(M, Y, P) \wedge setof(cf(X, M, P), seq(X, M, P), PredM) \wedge update(S_1, PredM, \{cf(M, Y, P)\}, S_2)$

(P2) $exists_non_executed_pred(M, P, S_1) \leftarrow seq(X, M, P) \wedge holds(not(cf(X, M, P)), S_1)$

expected to hold a value before the activity is executed. Second, an activity may produce a particular value (*output* relation), causing the item to get a new value. If it has no value yet, it is created, otherwise it is overwritten. It is worth noting that the item flow is not necessarily imposed over the control flow, but they interact for the definition of the process behavior. For instance, an activity expecting a value from a given item, may also cause a deadlock if this condition is never satisfied.

The item flow is modeled through the fluent $wrtn(A, It, P)$ (*wrtn* stands for "written") representing the situation in which the item It has been produced by the activity A in the enactment of the process P. In order to handle item manipulation, the semantics of task enactment (Rules A1, A2) is extended as follows:

$result(S_1, begin(A), S_2) \leftarrow task(A) \wedge holds(cf(X, A, P), S_1) \wedge \neg blocked_input(A, P, S_1) \wedge$
$input(A, It, P) \wedge update(S_1, \{cf(X, A, P)\}, \{en(A, P)\}, S_2)$

$result(S_1, complete(A), S_2) \leftarrow task(A) \wedge holds(en(A, P), S_1) \wedge seq(A, Y, P) \wedge$
$setof(wrtn(A, It, P), output(A, It, P), U) \wedge update(S_1, \{en(A, P)\}, \{cf(A, Y, P)\} \cup U, S_2)$

where $blocked_input(A, P, S_1)$ holds if, at a state S_1 during the enactment of process P, there exists some input item It for A that has not been produced. Thus,

$blocked_input(A, P, S_1) \leftarrow input(A, It, P) \wedge \neg updated_item(It, P, S_1)$
$updated_item(It, P, S_1) \leftarrow holds(wrtn(B, It, P), S_1))$

The case of compound activities can be treated in a similar way and is omitted.

4 Semantic Annotation

In the previous section we have shown how a procedural representation of a BPS can be modeled in our rule-based framework as an activity workflow. However, not all the relevant knowledge regarding process enactment is captured by a workflow model, which defines the planned order of operations but does not provide an explicit representation of the domain knowledge regarding the entities involved in such a process, i.e., the business environment in which processes are carried out.

Similarly to proposals like *Semantic BPM* [32] and *Semantic Web Services* [22], we will make use of *semantic annotations* to enrich the procedural knowledge specified by a BPS with domain knowledge expressed in terms of a given business reference ontology. Annotations provide two kinds of ontology-based information: *(i)* formal definitions of the basic entities involved in a process (e.g., activities, actors, items) to specify their meaning in an unambiguous way (*terminological* annotations), and *(ii)* specifications of preconditions and effects of the enactment of flow elements (*functional* annotations).

4.1 Reference Ontology

A business reference ontology is intended to capture the semantics of a business scenario in terms of the relevant vocabulary plus a set of axioms (TBox) that define the intended meaning of the vocabulary terms. In order to represent the semantic annotations of a BPAL BPS in a uniform way, we will consider ontologies falling within the OWL 2 RL profile (See Section 2.2), and hence expressible as sets of rules. An OWL 2 RL ontology is represented as a set $\mathcal{O}$ of rules, consisting of a set of facts of the form $t(s, p, o)$, called *triples*, encoding the specific OWL TBox, along with the rules that are common to all OWL 2 RL ontologies, such as the ones of Table 2.

In Table 6 we show an example of business reference ontology for the annotation of the Handle Order process depicted in Figure 1. For the sake of conciseness and clarity, the axioms of ontology are represented as DL expressions, instead of sets of triples. The translation into triple form can be done automatically as shown in [28, 59].

Table 6: Business Reference Ontology excerpt

Actors	
$Organizational_Actor \sqsubseteq Actor$	$Human_Actor \sqsubseteq Actor$
$Corporate_Customer \sqsubseteq Organizational_Actor$	$Employee \sqsubseteq Human_Actor$
$Department \sqsubseteq Organizational_Actor$	$Business_Partner \sqsubseteq Organizational_Actor$
$Accounting_Dpt \sqsubseteq Department$	$Supply_Dpt \sqsubseteq Department$
$Order_Mgt_Dept \sqsubseteq Department$	$Warehouse_Mgt \sqsubseteq Department$
$Carrier \sqsubseteq Organizational_Actor$	$Courier \sqsubseteq Carrier \sqcap Business_Partner$
$Carrier_Dpt \sqsubseteq Carrier \sqcap Department$	
Objects	
$ClosedPO \sqsubseteq Purchase_Order$	$ApprovedPO \sqsubseteq Purchase_Order$
$CancelledPO \sqsubseteq ClosedPO$	$FulfilledPO \sqsubseteq ClosedPO$
$UnavailablePL \sqsubseteq Part_List$	$AvailablePL \sqsubseteq Part_List$
$payment \sqsubseteq related$	$\exists\ payment^{-} \sqsubseteq Invoice$
$CancelledPO \sqcap ApprovedPO \sqsubseteq \bot$	$UnavailablePL \sqcap AvailablePL \sqsubseteq \bot$
$ApprovedPO \sqcap \exists related.Invoice \sqsubseteq FulfilledPO$	$Order \sqcap \exists related.UnavailablePL \sqsubseteq CancelledPO$
Processes	
$AuthorizingProcedure \sqsubseteq Process$	$Transportation \sqsubseteq Process$
$Payment \sqsubseteq Process$	$Invoicing \sqsubseteq Process$
$Communication \sqsubseteq Process$	$Refuse \sqsubseteq Communication$
$Rejecting \sqsubseteq Authorizing_Procedure$	$Accepting \sqsubseteq Authorizing_Procedure$
Relations	
$member \sqsubseteq related$	$content \sqsubseteq related$
$destination \sqsubseteq related$	$\exists member \sqsubseteq Human_Actor$
$\exists member^{-} \sqsubseteq Organizational_Actor$	

4.2 Terminological Annotation

A *terminological annotation* associates elements of a BPS with concepts of a reference ontology, in order to describe the former in terms of a suitable conceptualization of the underlying business domain provided by the latter. This association is specified by a set of OWL assertions of the form $BpsEl : \exists termRef.Concept$, where:

- *BpsEl* is an element of a BPS;
- *Concept* is either *i)* a named concept defined in the ontology, e.g., *Purchase_Order*, or *ii)* a complex concept, defined by a class expression, e.g., $Rejecting \sqcap \exists content.Purchase_Order$;
- *termRef* is an OWL property name.

We do not assume that every BPS element is annotated, nor that every concept is the meaning associated with some BPS element. Furthermore, different BPS elements could be annotated with respect to the same concept, to provide an alignment of the different terminologies and conceptualizations used in different BP schemas. E.g., the activities *bill_client* and *issue_invoice* occurring in different processes may actually refer to the same notion, suitably defined in the ontology.

Example 1. Examples of annotations related to the Handle Order process (Figure 1) are listed below. The item *order* is annotated with the *Purchase_Order* concept, while the participant *shipper* with the concept *Carrier*, which can be either an internal *Department* or a *Business_Partner*. A *sales_clerk* is defined as an *Employee*, which is part of the *Order_Mgt_Dpt*. The task *delivering* is defined as a *Transportation* related to some sort of *Product*. Finally, *notify_rejection* represents a *Communication* with a *Corporate_Customer*, and in particular, a *Refuse* related to *Purchase_Order*.

$order : \exists termRef.Purchase_Order$
$shipper : \exists termRef.Carrier$
$sales_clerk : \exists termRef.(Employee \sqcap \exists member.Order_Mgt_Dpt)$
$delivering : \exists termRef.(Transportation \sqcap \exists related.Product)$
$notify_rejection : \exists termRef.(Refuse \sqcap \exists content.Purchase_Order \sqcap$
$\exists destination.Corporate_Customer)$

4.3 Functional Annotation

By using the ontology vocabulary and axioms, we define semantic annotations for modeling the behavior of individual process elements in terms of *preconditions* under which a flow element can be executed, and *effects* on the state of the world after its execution. Preconditions and effects, collectively called *functional annotations*, can be used, for instance, to model input/output relations of activities with business entities. Fluents can represent the *properties* of a business entity affected by the execution of an activity at a given time during the execution of the process. A precondition specifies the properties a business entity must posses when an activity is enabled to start, and an effect specifies the properties of a business entity after having completed an activity. These aspects are only partially supported by current BP modeling notations, such as BPMN, in terms of data objects representing information storage during the BP enactment.

Functional annotations are formulated by means of the following relations:

- $pre(A, C, P)$, which specifies a fluent expression C, called *enabling condition*, that must hold to execute an element A in the process P;
- $eff(A, Q, E^-, E^+, P)$, which specifies the set E^- of fluents, called *negative effects*, that do not hold after the execution of A and the set of fluents E^+, called *positive effects*, that hold after the execution of A in the process P. Q is a fluent expression that must hold to complete the activity A. We assume that E^- and E^+ are disjoint sets of fluents, and the variables occurring in them also occur in Q.
- $c_seq(G, B, Y, P)$, which models a *conditional sequence flow* used to select the set of successors of decision points. G is a *guard* associated to the exclusive or inclusive branch gateway B, i.e., a fluent expressions that must hold in order to enable the flow element Y, successor of B in the process P. We also have the rule $seq(B, Y, P) \leftarrow c_seq(G, B, Y, P)$.

The enabling conditions, the guards and the negative and positive effects occurring in functional annotations are fluent expressions built from fluents of the form $t_f(s, p, o)$, corresponding to the OWL statement $t(s, p, o)$, where we adopt the usual *rdf*, *rdfs*, and *owl* prefixes for names in the OWL vocabulary, and the *bro* prefix for names relative to our specific examples. We assume that the fluents appearing in functional annotations are

either of the form $t_f(a, \textit{rdf:type}, c)$, corresponding to the unary atom $c(a)$, or of the form $t_f(a, p, b)$, corresponding to the binary atom $p(a, b)$, where a and c are *individuals*, while c and p are concepts and properties, respectively, defined in the reference ontology $\mathcal{O}$. Thus, fluents correspond to assertions about individuals, i.e., assertions belonging to the ABox of the ontology, and hence the ABox may change during process enactment due to the effects specified by the functional annotations, while $\mathcal{O}$, providing the ontology definitions and axioms, i.e., the TBox of the ontology, does not change.

Let us now present an example of specification of functional annotations. In particular, our example shows nondeterministic effects, that is, a case where a flow element A is associated with more than one pair (E^-, E^+) of negative and positive effects.

Example 2. Consider again the *Handle Order* process shown in Figure 1. After the execution of *create_order*, a purchase order is issued. This order can be approved or canceled upon execution of the activities *accept_order* and *cancel_order*, respectively. Depending on the inventory capacity checked during the *check_inventory* task, the requisition of parts performed by an external supplier is performed (*parts_auction*). Once that all the order parts are available, the order can be fulfilled and an invoice is associated with the order. This behavior is specified by the functional annotations reported in Table 7.

Table 7: Functional annotation for the Handle Order process

Flow Element	**Enabling Condition** (pre)	**Effects** (eff)
create_order	*true*	Q: *true* E^+: $\{t_f(o, \textit{rdf:type}, \textit{bro:Purchase_Order})\}$
accept_order	$t_f(O, \textit{rdf:type}, \textit{bro:Purchase_Order})$	Q: $t_f(O, \textit{rdf:type}, \textit{bro:Purchase_Order})$ E^+: $\{t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO})\}$
cancel_order	$t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO})$	Q: $t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO})$ E^-: $\{t_f(o, \textit{rdf:type}, \textit{bro:ApprovedPO})\}$ E^+: $\{t_f(o, \textit{rdf:type}, \textit{bro:CancelledPO})\}$
check_inventory	$t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO})$	Q: $t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO})$ E^+: $\{t_f(O, \textit{bro:related}, pl),$ $t_f(pl, \textit{rdf:type}, \textit{bro:Part_List})\}$
check_inventory	$t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO})$	
parts_auction	$t_f(PL, \textit{rdf:type}, \textit{bro:Part_List})$	Q: $t_f(PL, \textit{rdf:type}, \textit{bro:Part_List})$ E^+: $\{t_f(PL, \textit{rdf:type}, \textit{bro:AvailablePL})\}$
parts_auction	$t_f(PL, \textit{rdf:type}, \textit{bro:Part_List})$	Q: $and(t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO}),$ $t_f(PL, \textit{rdf:type}, \textit{bro:Part_List}))$ E^-: $\{t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO})\}$ E^+: $\{t_f(PL, \textit{rdf:type}, \textit{bro:UnavailablePL})\}$
bill_client	$t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO})$	Q: *true* E^+: $\{t_f(i, \textit{rdf:type}, \textit{bro:Invoice})\}$
payment	$and(t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO}),$ $t_f(I, \textit{rdf:type}, \textit{bro:Invoice}))$	Q: $and(t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO}),$ $t_f(I, \textit{rdf:type}, \textit{bro:Invoice}))$ E^+: $\{t_f(O, \textit{bro:payment}, I))\}$

Branch	**Successor**	**Guard**
g1	g3	$t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO})$
g1	g2	$not(t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO}))$
g3	parts_auction	$(t_f(PL, \textit{rdf:type}, \textit{bro:Part_List})$
g5	g2	$t_f(O, \textit{rdf:type}, \textit{bro:CancelledPO})$
g5	select_shipper	$t_f(O, \textit{rdf:type}, \textit{bro:ApprovedPO})$

4.3.1 Formal Semantics of Functional Annotations

In the presence of functional annotations, the enactment of a BPS is modeled by extending the *result* relation so as to take into account the *pre* and *eff* relations. We only consider the case of task execution. The other cases are similar and will be omitted.

Given a state S_1, a flow element A can be enacted if A is waiting for execution according to the control flow semantics, and its enabling condition C is satisfied, i.e., $holds(C, S_1)$ is true. Moreover, given an annotation $eff(A, Q, E^-, E^+, P)$, when A is completed in a given state S_1, then a new state S_2 is obtained by taking out from S_1 the set E^- of fluents and then adding the set E^+ of fluents. The execution of tasks considering functional annotations is then defined as:

$result(S_1, begin(A), S_2) \leftarrow task(A) \wedge holds(cf(B, A, P), S_1) \wedge pre(A, C, P) \wedge holds(C, S_1) \wedge update(S_1, \{cf(B, A, P)\}, \{en(A, P)\}, S_2)$

$result(S_1, complete(A), S_2) \leftarrow task(A) \wedge holds(en(A, P), S_1) \wedge eff(A, Q, E^-, E^+, P) \wedge holds(Q, S_1) \wedge seq(A, B, P) \wedge update(S_1, \{en(A, P)\} \cup E^-, \{cf(A, B, P)\} \cup E^+, S_2)$

Note that, since the variables occurring in E^+ and E^- are included in those of Q , the evaluation of $holds(Q, S_1)$ binds these variables to constants.

Similarly, the semantics of inclusive and exclusive branches is extended to evaluate the associated guard expressions, in order to determine the set of successors to be enabled. The execution of decision points is then defined as:

$result(S_1, complete(B), S_2) \leftarrow inc_branch(B) \wedge holds(cf(A, B, P), S_1) \wedge setof(cf(B, C, P), (c_seq(G, B, C, P) \wedge holds(G, S_1)), Succ) \wedge update(I, \{cf(A, B, P)\}, Succ, S_2)$

$result(S_1, complete(B), S_2) \leftarrow exc_branch(B) \wedge holds(cf(A, B, P), S_1) \wedge c_seq(G, B, C, P) \wedge holds(G, S_1) \wedge update(S_1, \{cf(A, B, P)\}, \{cf(B, C, P)\}, S_2)$

In order to evaluate a statement of the form $holds(t_f(s, p, o), X)$, where $t_f(s, p, o)$ is a fluent and X is a state, the definition of the *holds* predicate given previously must be extended to take into account the axioms belonging to the reference ontology $\mathcal{O}$. Indeed, we want that a fluent of the form $t_f(s, p, o)$ be true in state X not only if it belongs to X, but also if it can be inferred from the fluents in X and the axioms of the ontology.

For instance, let us consider the fluent $f = t_f(o, rdf{:}type, bro{:}CancelledPO)$. We can easily infer that f holds in a state that contains $\{t_f(o, rdf{:}type, bro{:}CancelledPO)\}$ (e.g., reachable after the execution of *cancel_order*) by using the rule $holds(F, X) \leftarrow F \in X$. However, by taking into account the ontology excerpt given in Table 6, we also want to be able to infer that f holds in a state that contains $\{t_f(o, rdf{:}type, bro{:}Purchase_Order)$, $t_f(o, bro{:}related, pl)$, $t_f(pl, rdf{:}type, bro{:}UnavailablePL)\}$ (e.g., a state reachable after the execution of *parts_auction*).

In our framework the inference of new fluents from fluents belonging to states is performed by including extra rules derived by translating the OWL 2 RL/RDF entailment rules as follows: every triple of the form $t(s, p, o)$, where s refers to an individual, is replaced by the atom $holds(t_f(s, p, o), X)$. Below we show exemplary rules (in particular, those required by our running example) for concept subsumption (1), role subsumption (2), domain restriction (3), transitive property (4), concept intersection[2] (5), existentially quantified formulae (6), and concept disjointness (7) . We refer the reader to [43] for the complete list of rules and the discussion of the OWL 2 RL rule-based semantics.

[2]Without loss of generality, unlike [43], we encode binary intersection instead of a general n-ary operator.

1. $holds(t_f(S, \textit{rdf:type}, C), X) \leftarrow holds(t_f(S, \textit{rdf:type}, B), X) \wedge t(B, \textit{rdfs:subClassOf}, C)$
2. $holds(t_f(S, P, O), X) \leftarrow holds(t_f(S, P1, O), X) \wedge t(P1, \textit{rdfs:subPropertyOf}, P)$
3. $holds(t_f(S, \textit{rdf:type}, C), X) \leftarrow holds(t_f(S, P, O), X) \wedge t(P, \textit{rdfs:domain}, C)$
4. $holds(t_f(S, P, O), X) \leftarrow holds(t_f(S, P, O_1), X) \wedge holds(t_f(O_1, P, O), X) \wedge$
   $t(P, \textit{rdf:type}, \textit{owl:TransitiveProperty})$
5. $holds(t_f(S, \textit{rdf:type}, C), X) \leftarrow t(C, \textit{owl:intersectionOf}, (C_1, C_2)) \wedge$
   $holds(t_f(S, \textit{rdf:type}, C_1), X) \wedge holds(t_f(S, \textit{rdf:type}, C_2), X)$
6. $holds(t_f(S, \textit{rdf:type}, C), X) \leftarrow t(S, \textit{owl:someValuesFrom,R}) \wedge t(R, \textit{owl:onProperty}, P) \wedge$
   $holds(and(t_f(S, P, I), t_f(I, \textit{rdf:type}, R)), X)$
7. $holds(\textit{false}, X) \leftarrow holds(t_f(I_1, \textit{rdf:type}, A), X) \wedge holds(t_f(I_2, \textit{rdf:type}, B), X) \wedge$
   $t(A, \textit{owl:disjointWith}, B)$

where *false* is a term representing $\bot$.

We denote by $\mathcal{A}$ the set of rules that encode the terminological and functional annotations, that is, (1) the OWL assertions of the form $BpsEl : \exists termRef.Concept$; (2) the facts defining the relations $pre(A, C, P)$, $\textit{eff}(A, Q, E^-, E^+, P)$, $c_seq(G, B, Y, P)$; (3) the rules for evaluating $holds(t_f(s, p, o), X)$ atoms (such as rules 1–7 above).

4.3.2 Change, Ramification and Consistency

The logical formalization of activity preconditions and effects given above has to be compared with various solutions to the *Frame* and *Ramification* problems proposed by the various AI formalisms for representing action and change.

The Frame Problem was formulated in [40] as the problem of "expressing a dynamical domain in logic without explicitly specifying which conditions are not affected by an action". Basically, it is concerned with *representational* issues, related to the effort needed to specify non-effects of actions, and *inferential* issues, related to the effort needed to actually compute these non-effects.

The Fluent Calculus addresses both the representational and inferential aspects of the Frame Problem [62] by modeling change as the difference between two states, caused by actions that deterministically result in a bounded number of direct (positive and negative) effects. These effects are captured by *state update axioms* specifying the fluents that are added or removed from a state. The rules defining the *result* relation introduced in Section 3.2 can be viewed as a specialized form of state update axioms.

The Ramification Problem [26] is the problem of representing and inferring information about indirect effects of actions. Indirect effects are not explicitly represented in action specifications, but follow from general laws (domain axioms) describing dependencies among fluents. In our framework, general laws are specified in the reference ontology TBox, whose axioms, as discussed in the previous section, are used in the derivation of additional t_f fluents from those belonging to a given state. Indirect effects may lead to undesired consequences when performing state update. For instance, let us consider the fluent $f = t_f(o, \textit{rdf:type}, \textit{bro:FulfilledPO})$. If we consider the ontology $\mathcal{O}$ given in Table 6, we can infer that f holds in a state S which contains $t_f(o, \textit{rdf:type}, \textit{bro:Purchase_Order})$, $t_f(o, \textit{bro:related}, i)$, and $t_f(i, \textit{rdf:type}, \textit{bro:Invoice})$. Now, assume that the set of negative effects of the subsequent activity a includes the fluent f. Then, after the state update determined by a, f still holds, in contrast with the intended meaning of negative effects.

Many approaches have been proposed to handle such a situation. Some of them are based on the computation of all the possible states s_i caused by the execution of action a in state s, such that: *i)* they comply with the domain axioms and the negative effects of a,

ii) they differ minimally from s (see, e.g., the Possible Model Approach - PMA [67]). This approach introduces a nondeterministic behavior in the state update that appears to be in contrast with the strong prescriptive nature of procedural BP models. Considering again the example above, the execution of a according to the PMA would result in three states: $s-\{t_f(o, \textit{rdf:type}, \textit{bro:Order})\}$, $s-\{t_f(o, \textit{bro:related}, i)\}$, and $s-\{t_f(i, \textit{rdf:type}, \textit{bro:Invoice})\}$.

Another solution proposed in the context of the Fluent Calculus, is based on *causal propagations* regulated by *causal relationships* [60], which specify how indirect effects are derived from direct effects and domain axioms. Causal relationships are then included in the state update axioms and applied until a fix-point is reached. This approach requires an additional formalism for the definition of causal relationships, and the burden for users of providing additional domain-dependent assertions, which cannot be represented within the ontology.

Here we follow a different approach based on the following *consistency condition*, which has to be enforced by every reachable state of a BPS: (i) no contradiction can be derived from the fluents belonging to the state by using the state independent axioms of the reference ontology, and (ii) no negative effect of an activity holds after its execution. Formally, we say that *eff* is *consistent* with process P if, for every flow element A and states S_1, S_2, the following implication is true:

If S_1 is reachable from the initial state of P and the relations $\textit{result}(S_1, \textit{complete}(A), S_2)$ and $\textit{eff}(A, E^-, E^+, P)$ hold,
Then $\mathcal{O} \cup \mathcal{A} \cup \{\neg \textit{holds}(\textit{false}, S_2)\}$ is consistent and, for all $F \in E^-$, $\mathcal{O} \cup \mathcal{A} \cup \{\neg \textit{holds}(F, S_2)\}$ is consistent.

This condition takes into account that, since $\mathcal{O} \cup \mathcal{A}$ is a definite logic program, it only allows the derivation of *positive* indirect effects, and thus, for all $F \in E^+$, $\mathcal{O} \cup \mathcal{A} \cup \{\textit{holds}(F, S_2)\}$ is consistent. We will show in Section 6 how the consistency condition can be checked by using the rule-based temporal logic we will present in the next section.

From a pragmatic perspective, the modeler is asked to refine the annotation of a BPS until a consistent description of the effects is achieved, possibly disambiguating the situations where underspecified effects may lead to hidden flaws.

5 Temporal Reasoning

In order to provide a general verification mechanism for behavioral properties, in this section we propose a model checking methodology based on a formalization of the temporal logic CTL (*Computation Tree Logic*, see [15] for a comprehensive overview) as a set of rules. Model checking is a widely accepted technique for the formal verification of BP schemas, as their execution semantics is usually defined in terms of states and state transitions, and hence the use of temporal logics for the specification and verification of properties is a very natural choice [24, 37].

CTL is a propositional temporal logic introduced for reasoning about the behavior of reactive systems. The behavior is represented as the tree of states that the system can reach, and each path of this tree is called a *computation path*. CTL formulas are built from: the constants *true* and *false*; a given set *Elem* of *elementary properties*; the connectives: $\neg$ ('not') and $\wedge$ ('and'); the linear-time operators along a computation path: **G** ('globally' or 'always'), **F** ('finally' or 'sometimes'), **X** ('next-time'), and **U** ('until'); the quantifiers over computation paths: **A** ('for all paths') and **E** ('for some path'). The abstract syntax of CTL is defined as follows.

Definition 1 (CTL formulas)**.** *A CTL formula F has the following syntax:*

$F ::= e \mid true \mid false \mid \neg F \mid F_1 \wedge F_2 \mid \mathbf{EX}(F) \mid \mathbf{EU}(F_1, F_2) \mid \mathbf{EG}(F)$

where e belongs to a given set Elem of elementary properties*.*

Other operators can be defined in terms of the ones given in Definition 1, e.g., $\mathbf{EF}(F) \equiv \mathbf{EU}(true, F)$ and $\mathbf{AG}(F) \equiv \neg\mathbf{EF}(\neg F)$ [15].

Usually, the semantics of CTL formulas is defined by introducing a *Kripke structure* $\mathcal{K}$, which represents the state space and the state transition relation, and by defining the satisfaction relation $\mathcal{K}, s \models F$, which denotes that a formula F holds in a state s of $\mathcal{K}$ [15]. In order to verify temporal properties of the behavior of a BPS P, we define a Kripke structure associated with P. The states are defined as finite sets of ground fluents and the state transition relation is based on the immediate reachability relation r between states defined in Section 3.2. The Kripke structure and the satisfaction relation will be encoded by sets of rules, hence providing a uniform framework for reasoning about the ontological properties and the behavioral properties of business processes.

A Kripke structure is a four-tuple $\mathcal{K} = \langle \mathcal{S}, \mathcal{I}, \mathcal{R}, \mathcal{L} \rangle$ defined as follows.

1. $\mathcal{S}$ is the finite set of all states, where a state is a finite set of ground fluents.
2. $\mathcal{I}$ is the *initial state* of BPS P, encoded by the rule:

   $initial(I, P) \leftarrow bp(P, S, E) \wedge I = \{cf(start, S, P)\}$
3. $\mathcal{R}$ is the *transition relation*, which is defined as follows: $\mathcal{R}(X, Y)$ holds iff $r(X, Y)$ holds, where r is the predicate defined in Section 3.2, i.e., $\mathcal{R}(X, Y)$ holds iff there exists an action A that can be executed in state X leading to state Y.
4. $\mathcal{L}$ is the *labeling function*, which associates with each state X the set of fluents F such that $\mathcal{O} \cup \mathcal{A} \models holds(F, X)$.

In the definition of Kripke structure given in [15], the transition relation $\mathcal{R}$ is assumed to be *total*, that is, every state S_1 has at least one *successor state* S_2 for which $\mathcal{R}(S_1, S_2)$ holds. This assumption is justified by the fact that reactive systems can be thought as ever running processes. However, this assumption is not realistic in the case of business processes, for which there is always at least one state with no successors, namely one where the *end* event of a BPS has been completed. For this reason the semantics of the temporal operators given in [15], which refers to *infinite* paths of the Kripke structure, is suitably changed here, according to [3], by taking into consideration *maximal* paths, i.e., paths that are either infinite or end with a state that has no successors, called a *sink*.

Definition 2 (Maximal Path)**.** *A* maximal path *in $\mathcal{K}$ starting from a state S_0 is either*

- *an* infinite *sequence of states $S_0\, S_1 \ldots$ such that $S_i\, \mathcal{R}\, S_{i+1}$, for every $i \geq 0$; or*
- *a* finite *sequence of states $S_0\, S_1\, \ldots\, S_k$, with $k \geq 0$, such that:*
  1. *$S_i\, \mathcal{R}\, S_{i+1}$, for every $0 \leq i < k$, and*
  2. *there exists no state $S_{k+1} \in \mathcal{S}$ such that $S_k\, \mathcal{R}\, S_{k+1}$.*

The semantics of CTL operators can be encoded by extending the definition of the predicate *holds*. Below we list the semantics of those operators and the corresponding rule-based formalization.

$\mathbf{EX}(F)$ holds in state S_0 if F holds in a successor state of S_0:

$holds(ex(F), S_0) \leftarrow r(S_0, S_1) \wedge holds(F, S_1)$

$\mathbf{EU}(F_1, F_2)$ holds in state S_0 if there exists a maximal path π: $S_0\ S_1 \ldots$ such that for some S_n in π we have that F_2 holds in S_n and, for $j = 0, \ldots, n-1$, F_1 holds in S_j:

$holds(eu(F_1, F_2), S_0) \leftarrow holds(F_2, S_0)$
$holds(eu(F_1, F_2), S_0) \leftarrow holds(F_1, S_0) \wedge r(S_0, S_1) \wedge holds(eu(F_1, F_2), S_1)$

$\mathbf{EG}(F)$ holds in a state S_0 if there exists a maximal path π starting from S_0 such that F holds in each state of π. Since the set of states is finite, $\mathbf{EG}(F)$ holds in S_0 if there exists a finite path $S_0\ \ldots\ S_k$ such that, for $i = 0, \ldots, k$, F holds in S_i, and either (1) $S_j = S_k$, for some $0 \leq j < k$, or (2) S_k is a sink state. Thus, the semantics of the operator **EG** is encoded by the following rules:

$holds(eg(F), S_0) \leftarrow fpath(F, S_0, S_0)$
$holds(eg(F), S_0) \leftarrow holds(F, S_0) \wedge r(S_0, S_1) \wedge holds(eg(F), S_1)$
$holds(eg(F), S_0) \leftarrow sink(S_0) \wedge holds(F, S_0)$

where: (i) the predicate $fpath(F, X, X)$ holds if there exists a path from X to X itself, consisting of at least one r arc, such that F holds in every state on the path:

$fpath(F, X, Y) \leftarrow holds(F, X) \wedge r(X, Y)$
$fpath(F, X, Z) \leftarrow holds(F, X) \wedge r(X, Y) \wedge\ \ fpath(F, Y, Z)$

and (ii) the predicate $sink(X)$ holds if X has no successor state.

Finally, the following rules define the properties characterizing the initial and the final state of a process:

$holds(F, s_0(P)) \leftarrow initial(I, P) \wedge holds(F, I)$
$holds(final(P), X) \leftarrow bp(P, S, E) \wedge holds(cf(E, end, P), X)$

The rules defining the semantics of the operator **EG** are similar to the constraint logic programming definition proposed in [45]. However, as already mentioned, in this paper we refer to the notion of maximal path instead of infinite path. Similarly to [45], our definition of the semantics of **EG** avoids the introduction of greatest fixed points of operators on sets of states which are often required by the approach described in [15]. Indeed, the rules defining $holds(eg(F), S_0)$ are interpreted according to the usual least fixpoint semantics (i.e., the least Herbrand model [38]).

The encoding of the satisfaction relation for other CTL operators, e.g, **EF** and **AG**, follows from the equivalences defining them [15]. It is worth noting that in some special cases the assumption that paths are maximal, but not necessarily infinite, matters [3]. For instance, if S_0 is a sink state, then $holds(ag(F), S_0)$ is true iff $holds(F, S_0)$ is true, since the only maximal path starting from S_0 is the one constituted by S_0 only. Finally, we would like to note that the definition of the CTL semantics given here is equivalent to the one in [15] in the presence of infinite computation paths only.

6 Reasoning Services

Our rule-based framework supports several reasoning services that can combine complex knowledge about business processes from different perspectives, such as the workflow structure, the ontological description, and the behavioral semantics. In this section we will illustrate three such services: verification, querying, and trace compliance checking.

Let us consider the following sets of rules: (1) $\mathcal{B}$, representing a set of BP schemas and the BP meta-model defined in Section 3.1, (2) $\mathcal{T}$, defining the behavioral semantics

presented in Section 3.2, (3) $\mathcal{O}$, collecting the OWL triples and rules that represent the business reference ontology defined in Section 4.1, (4) $\mathcal{A}$, encoding the annotations defined in Sections 4.2 and 4.3, and (5) $\mathcal{CTL}$, defining the semantics of CTL presented in Section 5.

Let $\mathcal{KB}$ be the set of rules $\mathcal{B} \cup \mathcal{T} \cup \mathcal{O} \cup \mathcal{A} \cup \mathcal{CTL}$. $\mathcal{KB}$ is called a *Business Process Knowledge Base* (BPKB). It is straightforward to show that $\mathcal{KB}$ is stratified, and hence its semantics is unambiguously defined by its perfect model $Perf(\mathcal{KB})$ (see Section 2.3).

6.1 Verification

In the following we present some examples of properties that can be specified and verified in our framework. A property is specified by a predicate *prop* defined by a rule C in terms of the predicates defined in $\mathcal{KB}$. The verification task is performed by checking whether or not $prop \in Perf(\mathcal{KB} \cup \{C\})$.

(1) A very relevant behavioral property of a BP p is that from any reachable state, it is possible to complete the process, i.e., reach the final state. This property, also known as *option to complete* [64], can be specified by the following rule, stating that the property *opt_com* holds if the CTL property $\mathbf{AG}(\mathbf{EF}(final(p)))$ holds in the initial state of p:

$opt_com \leftarrow holds(ag(ef(final(p))), s_0(p))$

(2) Temporal queries allow us to verify the consistency condition for effects introduced in Section 4.3. In particular, given a BPS p, inconsistencies due to the violation of some integrity constraint defined in the ontology by rules of the form $\bot \leftarrow G$ (e.g., concept disjointness) can be verified by defining the *inconsistency* property as follows:

$inconsistency \leftarrow holds(ef(false), s_0(p))$

(3) Another relevant property of a BPS is *executability* [66], according to which no activity reached by the control flow should be unable to execute due to some unsatisfied enabling condition. In our framework we can specify non-executability by defining a predicate *n_exec* which holds if it can be reached a state where some activity A is waiting for execution but is not possible to start its enactment.

$n_exec \leftarrow holds(ef(and(cf(A1, A, p), not(ex(en(A, p))))), s_0(p)) \wedge activity(A)$

(4) Temporal queries can also be used for the verification of *compliance rules*, i.e., directives expressing internal policies and regulations aimed at specifying the way an enterprise operates. In our Handle Order example, one such compliance rule may be that every *order* is eventually *closed*. In order to verify whether this property holds or not, we can define a *noncompliance* property which holds if it is possible to reach the final state of the process where, for some O, it can be inferred that O is an *order* which is not *closed*. In our example *noncompliance* is satisfied, and thus the compliance rule is not enforced. In particular, if the exception attached to the *accept order* task is triggered, the enactment continues with the *notify rejection* task (due to the guards associated to $g1$), and the order is never *canceled* nor *fulfilled*.

$noncompliance \leftarrow holds(ef(and(t_f(O, rdf{:}type, bro{:}Purchase_Order),$
$and(not(t_f(O, rdf{:}type, bro{:}ClosedPO)), final(p))), s_0(ho))$

6.2 Retrieval

The inference mechanism based on SLG-resolution can be used for computing boolean answers to ground queries, but also for computing, via unification, substitutions for variables occurring in non-ground queries. By exploiting this query answering mechanism we

can easily provide, besides the verification service described in the previous section, also reasoning services for the retrieval of process fragments.

The following queries show how process fragments can be retrieved according to different criteria. For sake of readability, we introduce the relation $\sigma(A, C)$ as an abbreviations for the OWL expression $A : \exists termRef.C$ encoding terminological annotations.

Query q_1 computes every activity A performed by a *Carrier* and realizing a *Transportation* (e.g., *delivering*) ; q_2 computes every decision point (exclusive branch) G occurring along a path of a BPS P delimited by two activities A and B, where the former operates on *orders* (e.g., *create_order*) and the latter is included in the results of q_1; finally, q_3 retrieve all the activities operating on *orders* which precede (in every possible execution) a *Transportation* performed by a *Carrier* (e.g., *create_order*).

$q_1(A) \leftarrow activity(A) \wedge assigned(A, C, P) \wedge \sigma(C, bro{:}Carrier) \wedge \sigma(A, bro{:}Transportation)$

$q_2(A, G, B, P) \leftarrow q_1(B) \wedge output(A, I, P) \wedge \sigma(I, bro{:}Purchase_Order) \wedge reachable(A, G, P) \wedge reachable(G, B, P)$

$q_3(A, B, P) \leftarrow q_1(B) \wedge output(A, I, P) \wedge \sigma(I, bro{:}Purchase_Order) \wedge reachable(A, B, P) \wedge holds(not(eu(not(en(A, P)), en(B, P))), s_0(P))$

6.3 Trace Compliance

The execution of a process is modeled as an *execution trace* (corresponding to a plan in the Fluent Calculus), i.e., a sequence of actions of the form $[act(a_1), \ldots, act(a_n)]$ where *act* is either *begin* or *complete*. The predicate $trace(S_1, T, S_2)$ defined below holds if T is a sequence of actions that lead from state S_1 to state S_2:

$trace(S_1, [\,], S_2) \leftarrow S_1 = S_2$

$trace(S_1, [A|T], S_2) \leftarrow result(S_1, A, U) \wedge trace(U, T, S_2)$

A *correct trace* T of a BPS P is a trace that leads from the initial state to the final state of P, that is:

$ctrace(T, P) \leftarrow initial(I, P) \wedge trace(I, T, Z) \wedge holds(final(P), Z)$

Execution traces are commonly stored by BPM systems as process logs, representing the evolution of the BP instances that have been enacted. The correctness of a trace t with respect to a given BPS p can be verified by evaluating a query of the form $ctrace(t, p)$ where t is a ground list and p is a process name.

The rules defining the predicate *ctrace* can also be used to *generate* the correct traces of a process p that satisfy some given property. This task is performed by evaluating a query of the form $ctrace(T, p) \wedge cond(T)$, where T is a free variable and $cond(T)$ is a property that T must enforce. For instance, we may want to generate traces where the execution of a flow element a is followed by the execution of a flow element b:

$cond(T) \leftarrow concat(T_1, T_2, T) \wedge complete(a) \in T_1 \wedge complete(b) \in T_2$

7 Computational Properties

In this section we prove the soundness, completeness, and termination of query evaluation using SLG-resolution. We also provide an upper bound to the time complexity of query evaluation.

As mentioned in Section 2.3, the soundness and completeness of SLG-resolution with respect to the perfect model semantics is guaranteed for the class of queries that do not

flounder. In [38] a sufficient condition ensuring that a query does not flounder is based on the notion of *allowed* query and rule. In particular, a query is allowed if every variable occurring in it also occurs in one of its positive literals. Similarly, a rule is allowed if every variable occurring in it also occurs in a positive literal in its body. Unfortunately, not all rules in $\mathcal{KB}$ are allowed in the sense of [38]. For instance, the variables F and S occurring in the rule $holds(not(F), S) \leftarrow \neg holds(F, S)$, do not occur in any positive literal of the body.

We will now define a subclass of the allowed queries whose evaluation with respect to $\mathcal{KB}$ does not flounder. The definition of this subclass also takes into account the left-to-right selection strategy for literals. For any predicate defined in $\mathcal{KB}$, each argument denoting a state (i.e., a set of fluents) can be classified either as an *input argument* or as an *output argument*. This classification is often called a *mode* [1]. In particular, it can be shown that we can classify the arguments such that the following property holds: if a predicate is evaluated with all its input arguments bound to ground sets of fluents, then whenever the predicate succeeds all variables occurring in output arguments are bound to ground sets of fluents. For instance, for the predicate $result(S_1, A, S_2)$ the first argument is an input argument and the third argument is an output argument. For reasons of space we do not list here, for each predicate defined in $\mathcal{KB}$, the input or output classification of its arguments. The following notion is adapted from [1].

Definition 3. *A query $L_1 \wedge \ldots \wedge L_n$ is* well-moded *if, for $i = 1, \ldots, n$, every variable occurring in an input argument in L_i occurs in an output argument in L_j, for some $j \in \{1, \ldots, i-1\}$.*

Definition 4. *Let f be (a term representing) a CTL formula. A subformula e of f is* grounding *if e is a fluent and one of the following conditions hold: (i) f is a fluent and e is f, (ii) f is $and(f_1, f_2)$ and e is a grounding subformula of either f_1 or f_2, (iii) f is $ex(f_1)$ and e is a grounding subformula of f_1, (iv) f is $eu(f_1, f_2)$ (or, in particular, f is $ef(f_2)$) and e is a grounding subformula of f_2, (v) f is $eg(f_1)$ and e is a grounding subformula of f_1.*

Definition 5. *A query Q of the form $L_1 \wedge \ldots \wedge L_n$ is an* NF*-query (short for* Non-Floundering *query) if the following conditions hold:*

(1) Q is well-moded,

(2) for $i = 1, \ldots, n$, if L_i is of the form $holds(f, S)$, then all variables of f occur in fluents that are subformulas of f, and

(3) for each variable X of Q, the leftmost occurrence X_l of X in Q appears in a positive literal L_j, with $1 \leq j \leq n$, such that either (3.1) L_j has predicate different from 'holds' or (3.2) $L_j = holds(f, S)$ and X_l appears in a grounding subformula of f.

A rule of the form $A \leftarrow L_1 \wedge \ldots \wedge L_n$ is an NF-rule*, if the following conditions hold:*

(4) no variable ranging over states occurs in A,

(5) every variable occurring in A also occurs in $L_1 \wedge \ldots \wedge L_n$, and

(6) $L_1 \wedge \ldots \wedge L_n$ is an NF-query.

For example, the queries and rules presented in Sections 6.1 and 6.2 are all NF. The query $holds(eu(en(A, p), true), s_0(p)) \wedge \neg task(A)$ is not an NF-query, because $en(A, p)$ is not a grounding subformula of $eu(en(A, p), true)$. This query flounders, as the non-ground negative literal $\neg task(A)$ will be selected after the success of $holds(eu(en(A, p), true), s_0(p))$.

We assume that the query Q is defined by a single NF-rule $Q \leftarrow B$, where B is a conjunction of literals. The extension to the case where Q is defined by a set of NF-rules (like in Section 6.2) is straightforward.

Proposition 1. *Let $Q \leftarrow B$ be an NF-rule such that the predicate of Q does not occur in $\mathcal{KB}$. Then we have the following properties.*
(1) The query Q does not flounder with respect to $\mathcal{KB} \cup \{Q \leftarrow B\}$.
(2) Every answer for Q with respect to $\mathcal{KB} \cup \{Q \leftarrow B\}$ is a ground substitution for the variables in Q.

Proof. (Sketch) (1) Let us consider a one-step derivation $L \wedge Q_1 \xrightarrow{\theta} Q_2$. By cases on the form of L one can show that if $L \wedge Q_1$ is an NF-query, then Q_2 is an NF-query. Thus, by also using the fact that every ground atom is an NF-query, if $\neg A \wedge Q_n$ is a query generable from Q in any number of steps, then $\neg A \wedge Q_n$ is an NF-query. Therefore, all variables occurring in $\neg A$ must also occur in a positive literal to the left of $\neg A$, and hence $\neg A$ is a ground atom.

(2) Suppose, by contradiction, that an answer θ for Q is not a ground substitution. Let us consider the rule $Q \leftarrow B \wedge \neg R$, where R is any atom containing one of the variables that are not bound to a ground term in θ. $Q \leftarrow B \wedge \neg R$ is an NF-rule. We can construct a derivation from Q that eventually selects the non-ground literal $\neg R\theta$, and hence the query Q flounders with respect to $\mathcal{KB} \cup \{Q \leftarrow B \wedge \neg R\}$. □

Let us now show that the evaluation of every NF-query terminates by using SLG resolution. Given an atomic query Q, we define:

- Calls_Q as the least set of atoms satisfying the following properties:

  (1) $Q \in \mathit{Calls}_Q$;

  (2) if $A \in \mathit{Calls}_Q$ and either $A \xrightarrow{\theta_1} \ldots \xrightarrow{\theta_n} A' \wedge Q'$ or $A \xrightarrow{\theta_1} \ldots \xrightarrow{\theta_n} \neg A' \wedge Q'$, then $A' \in \mathit{Calls}_Q$;

- $\mathit{Answers}_Q$ as the set of atoms $A\theta$ such that $A \in \mathit{Calls}_Q$ and θ is an answer for A;

- Δ_Q as $\mathit{Calls}_Q \cup \mathit{Answers}_Q$.

The termination proof is based on the property that, for any query Q, Δ_Q is a finite set of atoms. This property is equivalent to the *bounded-term-size property* that in [14] has been shown to be a sufficient condition for termination of SLG-resolution [14].

Given a set S, by $|S|$ we denote the cardinality of S. Let P be a logic program, by Π_P we denote the maximum number of literals in the body of a rule in P. The following result is an adaptation of Theorem 5.4.3 in [14].

Theorem 1 (Termination of SLG-resolution)**.** *Let P be a logic program and Q be an atomic query. Suppose that there exists a finite set $\mathcal{D}$ of atoms such that $\Delta_Q \subseteq \mathcal{D}$. Then all answers for Q can be computed by SLG-resolution in $\mathcal{O}(|P| \times |\mathcal{D}|^{\Pi_P+1})$ steps.*

By applying Theorem 1 to the case where P is of the form $\mathcal{KB} \cup \{q(X) \leftarrow B\}$, we get the following result.

Proposition 2. *Suppose that $q(X) \leftarrow B$ is an NF-rule, where X is a tuple of $k \geq 0$ variables and the predicates of B are defined in $\mathcal{KB}$. Then, all answers for $q(X)$ can be computed by SLG-resolution in $\mathcal{O}(|\mathcal{KB}| \times (|\mathcal{F}|^k + (\|B\| \times |\mathcal{F}|^v \times |\mathcal{S}|) + |\mathcal{S}|^m)^{r+1})$ steps, where: (i) $\mathcal{F}$ is the set of ground fluents that can be defined in $\mathcal{KB}$, (ii) $\mathcal{S}$ is the set of possible states, that is, the powerset of $\mathcal{F}$, (iii) $\|B\|$ denotes the size (that is, the number of symbols) of B, (iv) v is the largest number of variables in a CTL formula in B, (v) m is the largest arity of a predicate in $\mathcal{KB}$, and (vi) r is the largest number of literals in the body of a rule in $\mathcal{KB} \cup \{q(X) \leftarrow B\}$.*

Proof. Suppose that Q is the query $q(X)$ defined by the NF-rule $q(X) \leftarrow B$, where X is tuple of $k \geq 0$ variables and the predicates of B are defined in $\mathcal{KB}$. Let us define the following set $\mathcal{D}$ of atoms, where $\mathcal{V}$ is a finite, sufficiently large set of variables, and $\mathcal{E}$ is the set of flow elements in $\mathcal{KB}$.

$\mathcal{D} = \{q(t) \mid t \in (\mathcal{E} \cup \mathcal{F} \cup \mathcal{V})^k\} \cup$

$\{holds(f, s) \mid f = f'\theta$, for some CTL-formula f' occurring as a subformula in a literal of B and substitution θ from variables to fluents, and $s \in \mathcal{S}\} \cup$

$\{p(u) \mid p\ (\neq holds)$ is an m-ary predicate defined in $\mathcal{KB}$ and $u \in (\mathcal{E} \cup \mathcal{F} \cup \mathcal{S} \cup \mathcal{V})^m\}$

Additionally, we assume that no two atoms in $\mathcal{D}$ are variants of each other.

We have that $|\mathcal{D}| \leq (|\mathcal{E}| + |\mathcal{F}| + 1)^k + (\|B\| \times |\mathcal{F}|^v \times |\mathcal{S}|) + (|\mathcal{E}| + |\mathcal{F}| + |\mathcal{S}| + 1)^m$. The fluents in $\mathcal{F}$ are defined by using the elements in $\mathcal{E}$, the constants from the ontology (which also occur in $\mathcal{KB}$), and the function symbols *cf*, *en*, t_f, and *wrtn*, and hence $|\mathcal{E}| \leq |\mathcal{F}|$. Moreover, $|\mathcal{S}| = 2^{|\mathcal{F}|}$. Thus, $|\mathcal{D}| \in \mathcal{O}(|\mathcal{F}|^k + (\|B\| \times |\mathcal{F}|^v \times |\mathcal{S}|) + |\mathcal{S}|^m)$. By Theorem 1, we get the thesis. □

By using Propositions 1 and 2, we get the following result.

Theorem 2 (Termination, Soundness, and Completeness of Query Evaluation in $\mathcal{KB}$)**.** *Let $Q \leftarrow B$ be an NF-rule such that the predicate of Q does not occur in $\mathcal{KB}$. Then:*
(1) the evaluation of Q with respect to $\mathcal{KB} \cup \{Q \leftarrow B\}$ using SLG-resolution terminates;
(2) Q succeeds with answer θ iff $Q\theta \in Perf(\mathcal{KB} \cup \{Q \leftarrow B\})$;
(3) for a ground rule of the form prop $\leftarrow$ holds(f, s), the evaluation of prop by using SLG-resolution terminates in polynomial time in $\|f\| \times |\mathcal{S}|$.

Proof. (1) The termination of query evaluation has been proved in Proposition 2.

(2) The soundness and completeness of query evaluation follows from Proposition 1 and from the soundness and completeness of SLG-resolution for non-floundering queries recalled in Section 2.3.

(3) If we consider the ground rule *prop* $\leftarrow$ *holds*(f, s), then in Proposition 2 we have $k = v = 0$. Since m and r do not depend on f or $\mathcal{S}$, we get the thesis. □

Proposition 2 above only gives a loose upper bound on the complexity of query evaluation. However, it is sufficient for showing that, in line with the complexity of the CTL verification problem [15], our verification method has polynomial running time with respect to the number of states that are potentially reachable during process enactment. Moreover, Theorem 2 shows that the use of OWL 2 RL elementary properties does not add more than polynomial complexity. A tighter complexity analysis could be done by directly analyzing the evaluation of queries with respect to $\mathcal{KB}$, instead of relying, as done above, on the general results provided by [14].

In practice, our fluent-based representation of the behavioral semantics determines a running time which is polynomial in the number of flow elements that are concurrently enacted plus the number of fluents that are added to states by functional annotations. Usually, this number is much smaller than the cardinality of the powerset of $\mathcal{F}$. Indeed, the experimental results reported in Section 8.2 show that verification and querying are feasible for medium sized, non-trivial processes.

The termination of trace correctness checking can be proved under assumptions similar to the ones of Theorem 2. However, stronger assumptions are needed for the termination of trace generation in the case where we want to compute the set of *all* correct traces satisfying a given condition, as this set may be infinite in the presence of cycles.

8 Implementation

In the following we describe the BPAL Platform, a prototypical implementation of the framework discussed so far (Section 8.1), and we then discuss an experimental evaluation of the reasoner performances (Section 8.2).

8.1 Tool Description

The BPAL platform[3] is implemented as an Eclipse Plug-in[4], whose main components are depicted in the functional view in Figure 2. It provides the *BPKB Editor* to assist the user through a graphical interface in the definition of a BPKB, and the *BPAL Reasoner*, based on an LP engine, able to operate on the BPKB through the query language *QuBPAL*, designed for interrogating a repository of semantically enriched BPs.

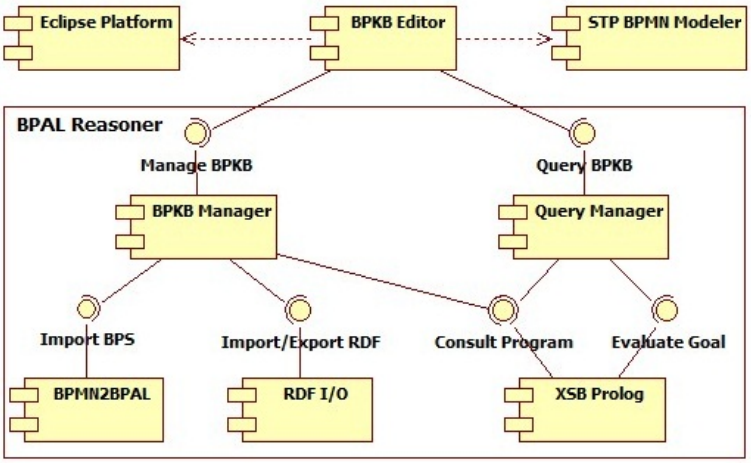

Figure 2: Functional view of the BPAL platform

8.1.1 Query Language

QuBPAL is an expressive query language for a BPKB based on the theoretical framework presented in this paper (a preliminary specification has been discussed in [55]). It does not require the user to understand the technicalities of the underlying LP platform,

[3] A video demonstration is available at `http://www.youtube.com/watch?v=xQkapzjhO7g`

[4] `http://www.eclipse.org/`

since QuBPAL queries are SELECT-WHERE statements intended to be automatically translated to logic programs, and then evaluated by using standard LP engines.

The SELECT statement defines the output of the query evaluation, which can be a boolean value, variables occurring in the WHERE statement, and a process selector representing either a BPS or a BPS fragment. The WHERE statement specifies an expression that restricts the set of data returned by the query, built from the set of the predicates defined in the BPKB (including CTL operators) and the connectives AND, OR, NOT, and the predicate = with the standard logic semantics. In the queries we use question mark to denote variables (e.g., *?x*), and we use the notation *?x::C* to indicate the terminological annotation of a variable, i.e., $x : \exists termRef.C$.

It is worth noting that the representation of OWL/RDFS resources as sets of triples, which directly encode the underlining RDF graph, allows us to pose queries over the ontology in a form very close to the SPARQL (SPARQL Protocol and RDF Query Language) standard [49], defined by the World Wide Web Consortium and widely accepted in the Semantic Web community. SPARQL is in fact designed to query RDF resources, that essentially are organized as directed and labeled graphs, by matching graph pattern over RDF graphs. Graph patterns are in turn specified as triples where variables can occur in every position (i.e., atoms of the form $t(a_1, a_2, a_3)$), along with their conjunctions and disjunctions. In this sense, while providing additional *primitives* to be used specifically for querying BPs, the ontology-related reasoning is specified in a QuBPAL query accordingly to consolidated Semantic Web standards.

To provide some insights about the language, we report in the following two examples of QuBPAL queries. The first one represents the formulation of the verification criteria for the compliance rule discussed at Point (4) of Section 6.1. The second one is the QuBPAL translation of the query q_3 discussed in Section 6.2.

```
SELECT <>
    WHERE [EF (final(ho) AND t(?o, rdf : type, bro : PurchaseOrder) AND NOT
    t(?o, rdf : type, bro : ClosedPO)| ho]
SELECT ?a ?p
    WHERE output(?a, ?i :: bro : Purchase_Order, ?p) AND reachable(?a, ?b, ?p) AND
    activity(?b :: bro : Transportation) AND assigned(?b, ?c :: bro : Carrier, ?p)
    AND [ NOT EU ( NOT en(?a, ?p), en(?b, ?p) ) | ?p]
```

8.1.2 BPKB Editor

This component provides a graphical user interface to define a BPKB and to interact with the BPAL Reasoner. A screen-shot of the main components of the GUI is depicted in Figure 3.

- The left panel (Figure 3.a) is the Package Explorer, providing a tree view of the resources available in the workspace, including BP schemas and ontologies.
- The central panel (Figure 3.b) is the BP Modeling View, based on the STP BPMN Modeler[5], comprising an editor and a set of tools to model business process diagrams using the BPMN notation.
- On the left (Figure 3.c), the Ontology View allows for the visualization of OWL ontologies, published on the Internet or locally stored.

[5] http://www.eclipse.org/soa

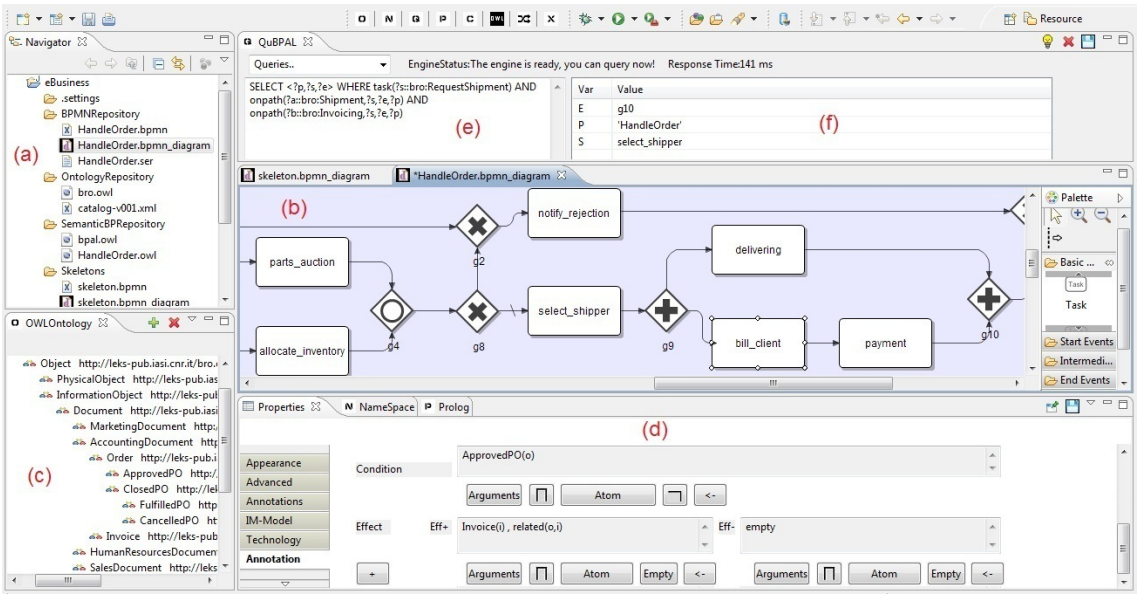

Figure 3: GUI of the BPAL platform

- The bottom panel (Figure 3.d) is the Annotation View, an editor for the annotation of process elements with respect to the reference ontology.
- The top-central panel (Figure 3.e) is the QuBPAL View, that provides a query prompt to access the BPAL reasoner through the query mechanism. Results can be consulted in the result panel (Figure 3.f).

8.1.3 BPAL Reasoner

This component implements the reasoning methods described in Section 6 by using the *XSB Prolog*[6] system [58], which is a Logic Programming system based on the SLG-resolution inference strategy recalled in Section 2.3. As proved in Section 7, the tabling mechanism guarantees the termination of query evaluation and the polynomial time (in the size of the state space) verification of CTL properties.

Process schemas are imported into the BPKB from BPMN process models via the *BPMN2BPAL* interface. In order to ease the sharing and re-use of semantic meta-data, semantic information used and produced during the annotation process (i.e., reference ontologies and semantic annotations) can be exported and imported from OWL/RDF files by means of the *RDF I/O* module. The underlying XSB Prolog implementation of the rule-based reasoner can deal with either RDF, RDFS or OWL 2 RL ontologies. The *BPKB Manager* handles the set-up and the interaction with the LP engine by initializing and updating a BPKB. After populating the BPKB, inference is essentially performed by posing queries to the *XSB Prolog* engine, connected through a Java/Prolog interface. To this end, the *Query Manager* exposes functionalities to translate QuBPAL queries into LP queries, evaluate them, and collect the results in a textual form or export them in an XML serialization.

[6] `http://xsb.sourceforge.net/`

8.2 Experimentation

The approach has been applied to real-world scenarios coming from end-users involved in the European Project BIVEE[7] and from the pilot conducted within a collaboration between the Italian CNR and SOGEI (ICT Company of the Italian Ministry of Finance). The former is related to the modeling of production processes in manufacturing oriented networked-enterprises, while the latter regards the procedural modeling of legislative decrees in the tax domain. The experiments we have conducted are encouraging and revealed the practical usability of the tool and its acceptance by business experts.

On a more technical side, the LP reasoner based on the XSB system shown a significant efficiency, since very sophisticated reasoning tasks have been performed on BPs of small-to-medium size (about one hundred of activities and several thousands of reachable states) in an acceptable amount of time and memory resources. Some empirical results are reported in the following, related to a dataset described in Table 8. We started by adapting a real world process, dealing with eProcurement, obtaining the BPS P, for which we report: the size, in terms of the number of flow elements; the number of reachable states; the number of exclusive, parallel, and inclusive gateways. As summarized in the table, the considered BPS does not contain logical errors (e.g., deadlocks) and is characterized by a considerable number of gateways, that is, branching/merging points (about 45% of the total number of elements). We then annotated in three different ways the process, obtaining P_1, P_2, P_3. For each one, in Table 8 we report: the number of reachable states; the coverage of the annotation, in terms of the percentage of the annotated flow elements; the average size of each state, in terms of the number of ontological assertions (i.e., t_f fluents) occurring in each state; the average size of the annotation, in terms of the number of t_f fluents occurring in the precondition/effect descriptions of the annotated flow elements; the errors exhibited by the BPS. In particular, P_1 has been annotated without preventing logical errors induced by the annotation, P_2 presents a revised version of P_1 annotation, further extended in P_3.

For the annotation of the BPS we adapted an ontology covering documents and production-related activities in the context of eProcurement and eBusiness, developed within the BIVEE project, comprising about 100 concepts.

Table 8: Annotated processes used in the evaluation

	Size	*States*	*XOR*	*PAR*	*OR*	*Errors*
P	87	821	14	14	6	No

	States	*Annotation Coverage*	*Average State Size*	*Average Annotation Size*	*Errors*
$\mathbf{P}_1$	944	35 %	7	3	2 non executable activities 150 inconsistent states 2 deadlocks
$\mathbf{P}_2$	2172	70 %	11	5	NO
$\mathbf{P}_3$	3866	100 %	16	8	NO

The experiments have been performed on an Intel laptop, with a 3 GHz Core 4 CPU, 8 GB RAM and Windows operating system. For each BPS we first tested the set-up of the reasoner, which include the translation of the BPKB into LP rules, their loading into the XSB reasoner, and the computation of the state space, i.e., the transitive closure of the *result* relation. Timing (measured in milliseconds) and memory occupation (measured in

[7]BIVEE: Business Innovation and Virtual Enterprise Environment (FoF-ICT-2011.7.3-285746)

megabytes) are reported in Table 9. We then run the queries presented in Section 6.1 and the last presented in Section 6.2, representing respectively: the verification of the option to complete (Q1), consistency condition (Q2), and executability (Q3) properties, an exemplary compliance rule (Q4) and a retrieval query (Q5). For each query, the average timing obtained in 10 runs is reported.

Table 9: Run-time phase evaluation

	State Space		***Query Evaluation***				
	Time	*Memory*	*Q1*	*Q2*	*Q3*	*Q4*	*Q5*
P	265	35	60	100	60	-	-
P_1	1030	210	110	2710	110	50	30
P_2	3300	670	530	4320	240	90	50
P_3	9720	1200	970	9250	405	105	60

To better understand the performed tests, additional considerations are needed. Firstly, the above CTL queries have been executed *after* the computation of the state space, which, due to the SLG-resolution strategy implemented by XSB, causes the population of the *tables* storing the intermediate results. The tables are then available in the subsequent queries, speeding up the computation. Secondly, to stress the engine, the evaluation of the performed queries requires the verification of ontology-based properties for each reachable state. Finally, the amount of required memory depends on the strategy adopted by the engine for the management of the tables. In the above experiments the default behavior has been adopted and, according to that, every intermediate result is materialized. This explains the large memory consumption, which, if needed, can be strongly reduced by introducing specific configurations to limit the use of tables, trading space for time.

It is also worth noting that no code optimization has been performed, since the executed Prolog program is the direct translation of the rules presented in this paper. Another remark regards the overhead introduced by the Java/Prolog bridge, which does not introduce a relevant performance degradation. Indeed, by running the same tests directly on XSB, without the Java infrastructure, the timings differ (up to a 10%) only in the presence of a large amount of results, mainly due to the inter-process data exchange.

9 Related Work

BP Modeling and Analysis. Among several mathematical formalisms proposed for defining a formal semantics of BP models, Petri nets [64] are the most used paradigm to capture the execution semantics of graph-based procedural languages (the BPMN case is discussed in [20]). Petri net models enable a large number of techniques for the control flow analysis of processes, but they do not provide a suitable basis to represent and reason about additional domain knowledge. In our framework we are able to capture the token game semantics underlying workflow models, and we can also declaratively represent constructs, such as exception handling behavior or synchronization of active branches only (inclusive merge), which, due to their non-local semantics, are cumbersome to capture in standard Petri nets. Furthermore, the logical grounding of our framework makes it easy to deal with the modeling of domain knowledge and the integration of reasoning services.

Program analysis and verification techniques have been largely applied to the analysis of process behavior, e.g., [24, 37]. These papers are based on the analysis of finite

state models through model checking techniques [15], where queries, formulated in some temporal logics, specify properties of process executions. However, these approaches are restricted to properties regarding the control flow only (e.g., properties of the ordering, presence, or absence of tasks in process executions), and severe limitations arise when ontology-related properties are included as part of the model to be checked.

Other approaches based on Logic Programming that are worth mentioning are [27, 53, 42]. [27] presents an approach to BP verification based on an extension of answer set programming with temporal logic and constraints, where the compliance of business rules is checked by bounded model checking techniques extended with constraint solving for dealing with conditions on numeric data. [53, 42] mainly focus on the analysis and on the enactment of flow models representing service choreographies, while we are not aware of specific extensions that deal with the semantic annotation of procedural process models with respect to domain ontologies.

Semantic Verification of BPs. There is a growing body of contributions beyond pure control flow verification [66, 25, 41, 19]. In [66] the authors introduce the notion of Semantic Business Process Validation, which aims at verifying properties related to the absence of logical errors which extend the notion of workflow soundness [64]. Validation is based on an execution semantics where token passing control flow is combined with the AI notion of state change induced by domain-related logical preconditions/effects. The main result is constituted by a validation algorithm that runs in polynomial time in the size of the workflow graph, under some restrictions on its structure and on the expressivity of the logic underlying the domain axiomatization, i.e., binary Horn clauses. This approach is focused on providing efficient techniques for the verification of specific properties, while the verification of arbitrary behavioral properties, such as the CTL formulae allowed in our framework, is not addressed. Moreover, our language for annotations, encompassing OWL 2 RL, is more expressive than binary Horn clauses. BP analysis techniques based on logical descriptions of effects of task execution are also proposed in [25, 41], but they introduce algorithms in an informal way, since a formal execution semantics is not provided, and a background ontology is not considered.

In [19] the authors discuss a CTL model checking method for annotated state transition systems, encoding the procedural behavior of Web Services interactions. Given a query, in the form of a CTL formula containing conjunctive subqueries, a *boolean* answer is computed in two steps: (1) a ground transition system is produced where each state contains all and only the description logic assertions *relevant* to the input query; (2) the grounded model is checked by a traditional propositional model checking algorithm. In contrast to our approach, the generation of the annotated transition system from a workflow model is neglected, and thus a semantics for activity preconditions/effects dealing also with the problems related to the state update is not given. Furthermore, our framework allows much more expressive reasoning services, since it is not limited to the boolean verification of CTL queries. On the technical side, our approach avoids the burden of integrating several tools, since both the temporal and ontological reasoning are performed by the LP inference engine. One relevant advantage of the LP translation is the possibility of computing answers according to a pure top-down, goal-oriented strategy, which avoids the need of preliminary grounding the model and possibly performing a large number of inferences that are not necessary for answering a given query.

Finally, we would like to mention a related research area, dealing with the verification of temporal properties in databases that evolve over time due to execution of actions operating on data (see [13] for a survey). Recently, [30] proposed *Knowledge and Action*

Bases (KABs), where actions, encoded as condition/action rules, modify the ABox of an ontology, encoded in a variant of the OWL 2 QL language. Under suitable restrictions, properties of KABs specified in the μ-calculus are shown to be decidable, and their verification can be reduced to finite-state model checking. KABs describe systems that may reach an infinite number of states, unlike our setting, where data are partially abstracted away[8], hence enforcing the reachable states to be a finite set. However, our framework is expressive enough to capture complex workflow specifications enriched with fluent expressions stated in terms of a background OWL 2 RL ontology. While the main goal of [30] is to provide theoretical results that characterize the decidability and (very high) complexity of KAB reasoning, our objective is more pragmatic and our formalization enables the implementation, through standard LP engines, of a wider set of (polynomial time) reasoning services, besides the verification of temporal properties.

Process Ontologies. The Process Specification Language (PSL) [9] is an ontology designed to formalize reasoning about processes in first-order logic. The basic structure that characterizes the behavior of a process in PSL is the *occurrence tree* (whose model is inspired by the Situation Calculus [51]), which contains all (infinite) sequences of occurrences of atomic activities starting from an initial state. Many extensions of PSL have been proposed to deal with time points, objects, agents, and resources. Although PSL is defined in first-order logic, which in principle makes behavioral specifications in PSL amenable to automated reasoning, it is mostly intended as a means to facilitate correct and complete exchange of process information among manufacturing systems, rather then for computation. Indeed, it is a very expressive framework whose associated reasoning tasks are intractable even for simple definitions, and undecidable in general, due to the adoption of unrestricted first-order logic. Furthermore, the systematic translation of procedural workflow descriptions into PSL has not been addressed, hence limiting its usability.

Several papers proposed the extension to BP management of techniques developed in the context of the Semantic Web[9]. To this end several meta-model process ontologies have been proposed, with the aim of specifying in a declarative, formal, and explicit way the modeling constructs, and enabling the use of domain ontologies for the semantic reconciliation of model contents. Some of them are derived from BP modeling notations (e.g., BPMN [17]), EPC [63], XPDL [29], Petri nets [11], while others have been designed in the context of interoperability, to overcome heterogeneities deriving from the adoption of different languages by mapping them to one common process ontology (e.g., GPO [36], BPMO [21]). The above approaches share some common features and goals: (1) they are based on standardized Web ontology languages; (2) they allow a machine-processable representation of BP models; (3) they enable query and search facilities; (4) they provide the means for relating BP models to existing business dictionaries and background knowledge. While a BPAL BPKB provides all the above features, supporting OWL 2 RL for ontological modeling, it also integrates behavioral modeling and a more expressive verification mechanism.

Semantic Web Services. Another stream of related papers regards the semantic enrichment of Web Services, where relevant work has been done within the OWL-S [12] and WSMO [52] initiatives. Both make an essential use of ontologies in order to facilitate the automation of discovering, combining and invoking electronic services over the Web. To

[8] basically, the set of individuals in the ontology is bounded and fixed a-priori; new values cannot be introduced during the enactment (e.g., by function terms)

[9] See the work conducted within the SUPER project: `http://www.ip-super.org/`

this end they describe services from two perspectives: from a *functional* perspective a service is described in terms of its functionality, preconditions and effects, input and output; from a process *perspective*, the service behavior is modeled as an orchestration of other services. However, in the above approaches the behavioral aspects are abstracted away, thus hampering the availability of reasoning services related to the execution of BPs. To overcome such limitations, several solutions for the representation of service compositions propose to translate the relevant aspects of the aforementioned service ontologies into a more expressive language, such as first-order logic. Among them, [57] adopts the high-level agent programming language Golog [51], [6, 44] rely on Situation Calculus variants, while [8, 5] are based on a direct translation of OWL-based service description into a Fluent Calculus theory. However, such approaches are mainly tailored to automated service composition (i.e., finding a sequence of service invocations such that a given goal is satisfied). Thus, the support provided for process definition, in terms of workflow constructs, is very limited and they lack a clear mapping from standard modeling notations. Furthermore, the adoption of a state-independent domain axiomatization (i.e., a DL TBOX) is not considered in the aforementioned approaches. In contrast, our framework allows a much richer procedural description of processes, directly corresponding to BPMN diagrams. Moreover, a reference ontology can be used to "enrich" process descriptions by means of annotations written in OWL 2 RL, one of the most widespread languages for ontology representation.

10 Conclusions

Summary

In this paper we discussed a methodological framework and a technical solution for the semantic enrichment of BP models, based on the synergic use of BPAL, a rule-based language adopted to provide a declarative representation of the procedural knowledge of a BP, and business ontologies, to capture the semantics of a business scenario. The resulting knowledge base provides a uniform and formal representation framework, suited for automated reasoning and equipped with a powerful inference mechanism supported by the programming systems developed in the area of Logic Programming.

BPAL is a rule-based formalism for modeling the structure and the behavior of a business process represented accordingly to a workflow perspective. It is essentially a process ontology, which provides a vocabulary, derived from BPMN, for specifying BPs, and an explicit description of its meta-model and execution semantics in terms of two core first-order logic theories which give formal definitions to the constructs of the language. In particular, from a control flow perspective, BPAL supports a relevant fragment of the BPMN standard, allowing us to deal with a large class of process models.

We then proposed an approach for the semantic enrichment of BPs, where BPAL BP schemas are related through a semantic annotation to a conceptualization of the business scenario formalized in a computational ontology. By integrating the rule-based ontology language OWL 2 RL with the structural and behavioral specification provided by BPAL, we are able to define a Business Process Knowledge Base (BPKB), as a collection of logical theories that provide a declarative representation of a repository of semantically enriched BPs.

On top of this knowledge representation framework, we built a number of reasoning services which allow the user to formulate complex queries that combine properties related to the structure, the behavioral semantics, and the ontological description of the BPs. We

showed how advanced resolution strategies, such as the tabled resolution implemented in the XSB Logic Programming system, guarantee a terminating, sound, and complete evaluation of the queries that can be issued over a BPKB.

Discussion

The rule-based approach followed in our framework offers several advantages. First of all, it enables the combination of the procedural and ontological perspectives in a very smooth and natural way, thus providing a uniform framework for reasoning on properties that depend on the sequence of operations that occur during process enactment and also on the domain where the process operates. Another advantage is the generality of the approach, which is open to further extensions, since other knowledge representation applications can easily be integrated, by providing a suitable translation to Logic Programming rules.

Furthermore, our approach does not introduce a new business process modeling language, but provides a framework where one can map and integrate knowledge represented by means of existing formalisms. This is very important from a pragmatic point of view, as one can express process-related knowledge by using standard modeling languages, while adding extra reasoning services. We have adopted BPMN as a graphical modeling notation, and its XML linear form to import and manipulate BP models, possibly designed through external BP Management Systems. For what concerns the ontology representation, we have committed to OWL, the current de-facto standard for ontology modeling and meta-data exchange. In essence, we have proposed a progressive approach, tailored to enhanced adaptability, where a business expert can start with the (commercial) tool and notation of his/her choice, and then enrich its functionalities with the formal framework we provide.

Finally, since our rule-based representation can be directly mapped to a class of logic programs, we can use standard Logic Programming systems to perform reasoning tasks such as verification and querying through a goal-oriented, efficient sound and complete evaluation procedure.

There are two main assumptions related to the practical applicability of our approach: the availability of ontologies and the willingness of an organization to describe their processes with semantic information. Clearly, enabling additional reasoning services comes at the price of additional modeling efforts, which may seriously hamper the adoption of our solution; this is a problem shared by many approaches based on Knowledge Representation techniques, in the Semantic Web related-research in particular. We now briefly discuss the impact of the above issues on the proposed approach.

The development of an ontology is a very complex task that requires the expertise of knowledge engineers and domain experts, and hence, high costs. Nevertheless, industrial products and services categorization standards, such as RosettaNet (`http://www.rosettanet.org/`) or eClass (`http://www.eclass-online.com/`), and libraries of standard business documents, such as UBL (`http://ubl.xml.org/`), reflect some degree of community consensus, and can thus be valuable input for creating business domain ontologies [31]. Also the growing interest for the publication of open data and their organization according to the Linked Data paradigm[10] increase the availability of publicly accessible terminological resources. Moreover, emerging methodologies for collaborative ontology building may be adopted here to lift existing resources (e.g., glossaries, organizational and data models) into formal theories [39]. That said, it should be noticed that

[10] http://linkeddata.org/

our framework does not require a heavy-weight, richly axiomatized ontology to work. The query capabilities can be still exploited even in the presence of a thesaurus only, which defines a set of terms whose meaning is agreed upon, possibly arranged in hierarchical structures. In this case, the annotation is reduced to tags taken from such a common glossary, but still retrieval and verification tasks with a practical relevance can be performed.

Also the semantic annotation is a time-consuming and error-prone task, which does not pay off if a small number of BPs has to be managed. However, in situations where hundreds of process models are available within an organization, and many collaborations with other departments or companies take place, the alignment of the adopted terminology and the reasoning facilities enabled by the semantic annotation may create a significant added-value[11]. Furthermore, once the ontologies are available, the effort required to the user for creating annotations amounts to browsing and selecting ontology concepts (see Section 8.1). In addition, we do not require that every BP is fully annotated; in many situations only parts of the model may be of interest for specific querying or verification tasks. Finally, approaches based on information retrieval and linguistic analysis can also be applied to support the annotation, suggesting correspondences between activity labels and terms defined in an ontology [18].

Future Work

The results presented in this paper leave several directions open for future research. First of all, we plan to push forward the empirical investigation of the impact of our proposal in each application scenario we are addressing, as reported in Section 8.2.

On the technical level, a relevant aspect to be further elaborated regards the adoption of query optimization techniques to enhance the reasoning approach. As it stands, the reasoner performs only simple optimizations based on the re-ordering of literals, and all the queries are evaluated with a pure goal-oriented, top-down approach, without any pre-processing of the knowledge base. We are confident that the query evaluation process can be strongly improved through more sophisticated query rewriting and program transformation techniques [48], which have been largely investigated in the area of Logic Programming.

We are also interested in applying the proposed framework in other phases of the BP life-cycle. In particular, the trace semantics of BPAL appears a suitable starting point to support: (i) querying at run-time, i.e., performed over a running instance of the process during its enactment; (ii) a-posteriori, i.e., over the execution logs of completed enactments, by adopting Inductive Logic Programming techniques, such as the ones presented in [35]; (iii) verification techniques for BPs in the presence of data constraints, by following approaches based on Constraint Logic Programming such as, for instance, the one proposed in [23].

Finally, we plan to extend the framework to also represent the execution-level process knowledge, and support the transition between conceptual and executable processes from a service-oriented perspective. That is, given a conceptual process model, Web services available in a repository are selected and possibly orchestrated to implement the process activities. The query-based support to process composition discussed in [54] represents a first contribution in that direction.

[11]See, e.g., the EU projects SUPER (http://www.ip-super.org/), Plug-it (http://plug-it-project.eu), COIN (http://www.coin-ip.eu/) and BIVEE (http://www.bivee.eu/).

References

[1] K. R. Apt. *From Logic Programming to Prolog*. Prentice Hall, London, UK, 1997.

[2] K. R. Apt and R. N. Bol. Logic Programming and Negation: A Survey. *Journal of Logic Programming*, 19/20:9–71, 1994.

[3] A. Arnold. *Finite transition systems: semantics of communicating systems*, chapter Transition System Logics. Prentice Hall International (UK) Ltd., Hertfordshire, UK, UK, 1994.

[4] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.

[5] G. Baryannis and D. Plexousakis. WSSL: A Fluent Calculus-Based Language for Web Service Specifications. In *CAiSE 2013*, volume 7908 of *LNCS*, pages 256–271. Springer, 2013.

[6] S. Battle et al. Semantic Web Services Ontology. `www.daml.org/services/swsf/1.0/swso/`, 2005.

[7] T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American*, May 2001.

[8] A. Bhuvaneswari and G. R. Karpagam. Applying Fluent Calculus for Automated and Dynamic Semantic Web Service Composition. In *Proc. of the 1st Int. Conf. on Intelligent Semantic Web-Services and Applications*, pages 16:1–16:6. ACM, 2010.

[9] C. Bock and M. Gruninger. PSL: A Semantic Domain for Flow Models. *Software and Systems Modeling Journal*, 4:209–231, 2004.

[10] D. Brickley and R. V. Guha. RDF Vocabulary Description Language 1.0: RDF Schema. W3C Recommendation, `www.w3.org/TR/2004/REC-rdf-schema-20040210/`, 2004.

[11] S. Brockmans, M. Ehrig, A. Koschmider, A. Oberweis, and R. Studer. Semantic Alignment Of Business Processes. In *8th Int. Conf. on Enterprise Information Systems*, pages 191–196. INSTICC Press, 2006.

[12] M. Burstein et al. OWL-S: Semantic Markup for Web Services. W3C Member Submission, `http://www.w3.org/Submission/OWL-S/`, 2004.

[13] D. Calvanese, G. De Giacomo, and M. Montali. Foundations of data-aware process analysis: A database theory perspective. In *Proceedings of the 32nd Symposium on Principles of Database Systems*, PODS '13, pages 1–12. ACM, 2013.

[14] W. Chen and D. S. Warren. Tabled Evaluation with Delaying for General Logic Programs. *J. ACM*, 43:20–74, January 1996.

[15] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. The MIT Press, 1999.

[16] A. De Nicola, M. Missikoff, M. Proietti, and F. Smith. An Open Platform for Business Process Modeling and Verification. In *Database and Expert Systems Applications, 21st International Conference*, volume 6261 of *LNCS*, pages 76–90. Springer, 2010.

[17] C. Di Francescomarino, C. Ghidini, M. Rospocher, L. Serafini, and P. Tonella. Semantically-Aided Business Process Modeling. In *International Semantic Web Conference*, LNCS 5823, pages 114–129. Springer, 2009.

[18] C. Di Francescomarino and P. Tonella. Supporting Ontology-Based Semantic Annotation of Business Processes with Automated Suggestions. *IJISMD*, 1(2):59–84, 2010.

[19] I. Di Pietro, F. Pagliarecci, and L. Spalazzi. Model Checking Semantically Annotated Services. *IEEE Trans. Software Eng.*, 38(3):592–608, 2012.

[20] R. M. Dijkman, M. Dumas, and C. Ouyang. Semantics and Analysis of Business Process Models in BPMN. *Inf. Softw. Technol.*, 50:1281–1294, November 2008.

[21] M. Dimitrov, A. Simov, S. Stein, and M. Konstantinov. A BPMO Based Semantic Business Process Modelling Environment. In *Semantic Business Process Modeling*, volume 251 of *CEUR-WS Proceedings*. CEUR-WS.org, 2007.

[22] D. Fensel et al. *Enabling Semantic Web Services: The Web Service Modeling Ontology*. Springer, 2006.

[23] F. Fioravanti, A. Pettorossi, M. Proietti, and V. Senni. Generalization strategies for the verification of infinite state systems. *Theory and Practice of Logic Programming*, 13(2):175–199, 2013.

[24] X. Fu, T. Bultan, and J. Su. Analysis of Interacting BPEL Web Services. In *Proc. of the Int. Conf. on World Wide Web*, pages 621–630. ACM, 2004.

[25] A. Ghose and G. Koliadis. Auditing Business Process Compliance. In *Proc. of Int. Conf. on Service-Oriented Computing*, volume 4749 of *LNCS*, pages 169–180. Springer, 2007.

[26] M. L. Ginsberg and D. E. Smith. Reasoning about action II: The qualification problem. *Artificial Intelligence*, 35(3):311–342, 1988.

[27] L. Giordano, A. Martelli, M. Spiotta, and D. T. Dupré. Business process verification with constraint temporal answer set programming. *TPLP*, 13(4-5):641–655, 2013.

[28] B. N. Grosof, I. Horrocks, R. Volz, and S. Decker. Description Logic Programs: Combining Logic Programs with Description Logic. In *Proc. of the 12th Int. Conf. on World Wide Web*, WWW '03, pages 48–57. ACM, 2003.

[29] A. Haller, W. Gaaloul, and M. Marmolowski. Towards an XPDL Compliant Process Ontology. In *Proc. of the IEEE Congress on Services*, pages 83–86. IEEE Computer Society, 2008.

[30] B. B. Hariri, D. Calvanese, M. Montali, G. D. Giacomo, R. D. Masellis, and P. Felli. Description logic knowledge and action bases. *J. Artif. Intell. Res. (JAIR)*, 46:651–686, 2013.

[31] M. Hepp. Products and Services Ontologies: A Methodology for Deriving OWL Ontologies from Industrial Categorization Standards. *Int. J. Semantic Web Inf. Syst.*, 2(1):72–99, 2006.

[32] M. Hepp et al. Semantic Business Process Management: A Vision Towards Using Semantic Web Services for Business Process Management. In *Proc. of Int. Conf. on e-Business Engineering*. IEEE Computer Society, 2005.

[33] E. Kindler. On the Semantics of EPCs: Resolving the Vicious Circle. *Data Knowl. Eng.*, 56(1):23–40, 2006.

[34] G. Klyne and J. J. Carroll. Resource Description Framework (RDF): Concepts and Abstract Syntax. W3C Recommendation, `http://www.w3.org/TR/rdf-concepts/`, February 2004.

[35] E. Lamma, P. Mello, F. Riguzzi, and S. Storari. Applying Inductive Logic Programming to Process Mining. In *Inductive Logic Programming*, volume 4894 of *LNCS*, pages 132–146. Springer, 2007.

[36] Y. Lin. *Semantic Annotation for Process Models: Facilitating Process Knowledge Management via Semantic Interoperability*. PhD thesis, Norwegian University of Science and Technology, 2008.

[37] Y. Liu, S. Müller, and K. Xu. A Static Compliance-Checking Framework for Business Process Models. *IBM Syst. J.*, 46:335–361, April 2007.

[38] J. W. Lloyd. *Foundations of logic programming*. Springer-Verlag New York, Inc., 1987.

[39] V. Ludovici, F. Smith, and F. Taglino. Collaborative Ontology Building in Virtual Innovation Factories. In *2013 Int. Conf. on Collaboration Technologies and Systems*, pages 443–450. IEEE, 2013.

[40] J. McCarthy and P. J. Hayes. Some Philosophical Problems from the Standpoint of Artificial Intelligence. In *Machine Intelligence*, volume 4, pages 463–502, 1969.

[41] H. Meyer. On the Semantics of Service Compositions. In *Proc. of the 1st Int. Conf. on Web Reasoning and Rule Systems*, volume 4524 of *LNCS*, pages 31–42. Springer, 2007.

[42] M. Montali, M. Pesic, W. M. P. van der Aalst, F. Chesani, P. Mello, and S. Storari. Declarative specification and verification of service choreographies. *ACM Trans. Web*, 4(1):3:1–3:62, Jan. 2010.

[43] B. Motik et al. OWL 2 Web Ontology Language Profiles (Second Edition). W3C Recommendation, `http://www.w3.org/TR/owl2-profiles/`, 2012.

[44] S. Narayanan and S. McIlraith. Analysis and Simulation of Web Services. *Comp. Networks*, 42:675–693, August 2003.

[45] U. Nilsson and J. Lübcke. Constraint Logic Programming for Local and Symbolic Model-checking. In *Computational Logic*, LNAI 1861. Springer, 2000.

[46] OMG. Business Process Model and Notation. http://www.omg.org/spec/BPMN/2.0, 2011.

[47] M. Penker and H.-E. Eriksson. *Business Modeling With UML: Business Patterns at Work*. Wiley, Jan. 2000.

[48] A. Pettorossi and M. Proietti. Transformation of logic programs: Foundations and techniques. *J. Log. Program.*, 19/20:261–320, 1994.

[49] E. Prud'hommeaux and A. Seaborne. SPARQL Query Language for RDF. W3C Recommendation, `http://www.w3.org/TR/2007/WD-rdf-sparql-query-20070326/`, January 2008.

[50] T. C. Przymusinski. On the Declarative Semantics of Deductive Databases and Logic Programs. In *Foundations of Deductive Databases and Logic Programming*. Morgan Kaufmann Publishers, 1988.

[51] R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press, Sept. 2001.

[52] D. Roman, U. Keller, H. Lausen, J. de Bruijn, R. Lara, M. Stollberg, Polleres, C. Feier, C. Bussler, and D. Fensel. Web Service Modeling Ontology. *Applied Ontology*, 1(1):77–106, 2005.

[53] D. Roman and M. Kifer. Semantic Web Service Choreography: Contracting and Enactment. In *Proc. of International Semantic Web Conference*, LNCS 5318, pages 550–566, 2008.

[54] F. Smith and D. Bianchini. Semi-Automatic Process Composition via Semantics-Enabled Sub-Process Selection and Ranking. In *Enterprise Interoperability V, Shaping Enterprise Interoperability in the Future Internet, Proc. of I-ESA 2012*, pages 177–187. Springer, 2012.

[55] F. Smith, M. Missikoff, and M. Proietti. Ontology-Based Querying of Composite Services. In *Business System Management and Engineering*, volume 7350 of *LNCS*, pages 159–180. Springer, 2010.

[56] F. Smith and M. Proietti. Rule-based Behavioral Reasoning on Semantic Business Processes. In *ICAART (2)*, pages 130–143. SciTePress, 2013.

[57] S. Sohrabi, N. Prokoshyna, and S. A. McIlraith. Web Service Composition via the Customization of Golog Programs with User Preferences. In *Conceptual Modeling: Foundations and Applications*, pages 319–334. Springer, 2009.

[58] T. Swift and D. Warren. XSB: Extending the power of Prolog using tabling. 12(1-2):157–187, 2012.

[59] H. ter Horst. Extending the RDFS Entailment Lemma. In D. Plexousakis, S. McIlraith, and F. van Harmelen, editors, *Proc. of the 3rd Int. Semantic Web Conf.*, LNCS. Springer, 2004.

[60] M. Thielscher. Ramification and Causality. *Artificial Intelligence*, 89:317–364, 1997.

[61] M. Thielscher. Introduction to the Fluent Calculus. *Electron. Trans. Artif. Intell.*, 2:179–192, 1998.

[62] M. Thielscher. From Situation Calculus to Fluent Calculus: State Update Axioms as a Solution to the Inferential Frame Problem. *Artificial Intelligence*, 111(1-2):277–299, July 1999.

[63] O. Thomas and M. Fellmann. Semantic EPC: Enhancing Process Modeling Using Ontology Languages. In *Semantic Business Process Modeling*, volume 251 of *CEUR-WS Proceedings*. CEUR-WS.org, 2007.

[64] W. M. P. van der Aalst. The Application of Petri Nets to Workflow Management. *Journal of Circuits, Systems, and Computers*, 8(1):21–66, 1998.

[65] H. Völzer. A New Semantics for the Inclusive Converging Gateway in Safe Processes. In *Proc. of Int. Conf. on Business Process Management*, LNCS 6336, pages 294–309. Springer, 2010.

[66] I. Weber, J. Hoffmann, and J. Mendling. Beyond Soundness: on the Verification of Semantic Business Process Models. *Distrib. Parallel Databases*, 27:271–343, June 2010.

[67] M. Winslett. Reasoning about Action Using a Possible Models Approach. In *AAAI*, pages 89–93. The MIT Press, 1988.